Dell PowerConnect J-Series Ethernet Switch

Complete Software Guide for Junos OS

Published: 2010-11-10

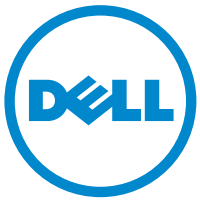

Dell
501 Dell Way
Round Rock , Texas 78682
United States
www.dell.com

Dell PowerConnect J-Series Ethernet Switch Complete Software Guide for Junos OS

Writing: Appumon Joseph, Aviva Garrett, Bhargava Y P, Brian Deutscher, Brooke Doverspike, Carolyn Harding, Greg Houde, Hareesh Kumar K N, Hemraj Rao S, Janet Bein, Katherine Kearns, Keldyn West, Praveen G R, Regina Roman, Shikha Kalra, Tim Harrington, Vinita Kurup
Editing: Cindy Martin, Rajan V K, Taffy Everts, Chanchal Agrawal
Illustration: Faith Bradford Brown
Cover Design: Christine Nay

Revision History
15 November 2010—Revision 2
4 June 2010—Revision 1

The information in this document is current as of the date listed in the revision history.

SOFTWARE LICENSE

The terms and conditions for using this software are described in the software license contained in the acknowledgment to your purchase order or, to the extent applicable, to any reseller agreement or end-user purchase agreement executed between you and Juniper Networks. By using this software, you indicate that you understand and agree to be bound by those terms and conditions.

Generally speaking, the software license restricts the manner in which you are permitted to use the software and may contain prohibitions against certain uses. The software license may state conditions under which the license is automatically terminated. You should consult the license for further details.

For complete product documentation, please see the Juniper Networks Web site at www.juniper.net/techpubs.

END USER LICENSE AGREEMENT

READ THIS END USER LICENSE AGREEMENT ("AGREEMENT") BEFORE DOWNLOADING, INSTALLING, OR USING THE SOFTWARE. BY DOWNLOADING, INSTALLING, OR USING THE SOFTWARE OR OTHERWISE EXPRESSING YOUR AGREEMENT TO THE TERMS CONTAINED HEREIN, YOU (AS CUSTOMER OR IF YOU ARE NOT THE CUSTOMER, AS A REPRESENTATIVE/AGENT AUTHORIZED TO BIND THE CUSTOMER) CONSENT TO BE BOUND BY THIS AGREEMENT. IF YOU DO NOT OR CANNOT AGREE TO THE TERMS CONTAINED HEREIN, THEN (A) DO NOT DOWNLOAD, INSTALL, OR USE THE SOFTWARE, AND (B) YOU MAY CONTACT JUNIPER NETWORKS REGARDING LICENSE TERMS.

1. **The Parties.** The parties to this Agreement are (i) Juniper Networks, Inc. (if the Customer's principal office is located in the Americas) or Juniper Networks (Cayman) Limited (if the Customer's principal office is located outside the Americas) (such applicable entity being referred to herein as "Juniper"), and (ii) the person or organization that originally purchased from Juniper or an authorized Juniper reseller the applicable license(s) for use of the Software ("Customer") (collectively, the "Parties").

2. **The Software.** In this Agreement, "Software" means the program modules and features of the Juniper or Juniper-supplied software, for which Customer has paid the applicable license or support fees to Juniper or an authorized Juniper reseller, or which was embedded by Juniper in equipment which Customer purchased from Juniper or an authorized Juniper reseller. "Software" also includes updates, upgrades and new releases of such software. "Embedded Software" means Software which Juniper has embedded in or loaded onto the Juniper equipment and any updates, upgrades, additions or replacements which are subsequently embedded in or loaded onto the equipment.

3. **License Grant.** Subject to payment of the applicable fees and the limitations and restrictions set forth herein, Juniper grants to Customer a non-exclusive and non-transferable license, without right to sublicense, to use the Software, in executable form only, subject to the following use restrictions:

a. Customer shall use Embedded Software solely as embedded in, and for execution on, Juniper equipment originally purchased by Customer from Juniper or an authorized Juniper reseller.

b. Customer shall use the Software on a single hardware chassis having a single processing unit, or as many chassis or processing units for which Customer has paid the applicable license fees; provided, however, with respect to the Steel-Belted Radius or Odyssey Access Client software only, Customer shall use such Software on a single computer containing a single physical random access memory space and containing any number of processors. Use of the Steel-Belted Radius or IMS AAA software on multiple computers or virtual machines (e.g., Solaris zones) requires multiple licenses, regardless of whether such computers or virtualizations are physically contained on a single chassis.

c. Product purchase documents, paper or electronic user documentation, and/or the particular licenses purchased by Customer may specify limits to Customer's use of the Software. Such limits may restrict use to a maximum number of seats, registered endpoints, concurrent users, sessions, calls, connections, subscribers, clusters, nodes, realms, devices, links, ports or transactions, or require the purchase of separate licenses to use particular features, functionalities, services, applications, operations, or capabilities, or provide throughput, performance, configuration, bandwidth, interface, processing, temporal, or geographical limits. In addition, such limits may restrict the use of the Software to managing certain kinds of networks or require the Software to be used only in conjunction with other specific Software. Customer's use of the Software shall be subject to all such limitations and purchase of all applicable licenses.

d. For any trial copy of the Software, Customer's right to use the Software expires 30 days after download, installation or use of the Software. Customer may operate the Software after the 30-day trial period only if Customer pays for a license to do so. Customer may not extend or create an additional trial period by re-installing the Software after the 30-day trial period.

e. The Global Enterprise Edition of the Steel-Belted Radius software may be used by Customer only to manage access to Customer's enterprise network. Specifically, service provider customers are expressly prohibited from using the Global Enterprise Edition of the Steel-Belted Radius software to support any commercial network access services.

The foregoing license is not transferable or assignable by Customer. No license is granted herein to any user who did not originally purchase the applicable license(s) for the Software from Juniper or an authorized Juniper reseller.

4. **Use Prohibitions.** Notwithstanding the foregoing, the license provided herein does not permit the Customer to, and Customer agrees not to and shall not: (a) modify, unbundle, reverse engineer, or create derivative works based on the Software; (b) make unauthorized

copies of the Software (except as necessary for backup purposes); (c) rent, sell, transfer, or grant any rights in and to any copy of the Software, in any form, to any third party; (d) remove any proprietary notices, labels, or marks on or in any copy of the Software or any product in which the Software is embedded; (e) distribute any copy of the Software to any third party, including as may be embedded in Juniper equipment sold in the secondhand market; (f) use any 'locked' or key-restricted feature, function, service, application, operation, or capability without first purchasing the applicable license(s) and obtaining a valid key from Juniper, even if such feature, function, service, application, operation, or capability is enabled without a key; (g) distribute any key for the Software provided by Juniper to any third party; (h) use the Software in any manner that extends or is broader than the uses purchased by Customer from Juniper or an authorized Juniper reseller; (i) use Embedded Software on non-Juniper equipment; (j) use Embedded Software (or make it available for use) on Juniper equipment that the Customer did not originally purchase from Juniper or an authorized Juniper reseller; (k) disclose the results of testing or benchmarking of the Software to any third party without the prior written consent of Juniper; or (l) use the Software in any manner other than as expressly provided herein.

5. **Audit.** Customer shall maintain accurate records as necessary to verify compliance with this Agreement. Upon request by Juniper, Customer shall furnish such records to Juniper and certify its compliance with this Agreement.

6. **Confidentiality.** The Parties agree that aspects of the Software and associated documentation are the confidential property of Juniper. As such, Customer shall exercise all reasonable commercial efforts to maintain the Software and associated documentation in confidence, which at a minimum includes restricting access to the Software to Customer employees and contractors having a need to use the Software for Customer's internal business purposes.

7. **Ownership.** Juniper and Juniper's licensors, respectively, retain ownership of all right, title, and interest (including copyright) in and to the Software, associated documentation, and all copies of the Software. Nothing in this Agreement constitutes a transfer or conveyance of any right, title, or interest in the Software or associated documentation, or a sale of the Software, associated documentation, or copies of the Software.

8. **Warranty, Limitation of Liability, Disclaimer of Warranty.** The warranty applicable to the Software shall be as set forth in the warranty statement that accompanies the Software (the "Warranty Statement"). Nothing in this Agreement shall give rise to any obligation to support the Software. Support services may be purchased separately. Any such support shall be governed by a separate, written support services agreement. TO THE MAXIMUM EXTENT PERMITTED BY LAW, JUNIPER SHALL NOT BE LIABLE FOR ANY LOST PROFITS, LOSS OF DATA, OR COSTS OR PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES, OR FOR ANY SPECIAL, INDIRECT, OR CONSEQUENTIAL DAMAGES ARISING OUT OF THIS AGREEMENT, THE SOFTWARE, OR ANY JUNIPER OR JUNIPER-SUPPLIED SOFTWARE. IN NO EVENT SHALL JUNIPER BE LIABLE FOR DAMAGES ARISING FROM UNAUTHORIZED OR IMPROPER USE OF ANY JUNIPER OR JUNIPER-SUPPLIED SOFTWARE. EXCEPT AS EXPRESSLY PROVIDED IN THE WARRANTY STATEMENT TO THE EXTENT PERMITTED BY LAW, JUNIPER DISCLAIMS ANY AND ALL WARRANTIES IN AND TO THE SOFTWARE (WHETHER EXPRESS, IMPLIED, STATUTORY, OR OTHERWISE), INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT. IN NO EVENT DOES JUNIPER WARRANT THAT THE SOFTWARE, OR ANY EQUIPMENT OR NETWORK RUNNING THE SOFTWARE, WILL OPERATE WITHOUT ERROR OR INTERRUPTION, OR WILL BE FREE OF VULNERABILITY TO INTRUSION OR ATTACK. In no event shall Juniper's or its suppliers' or licensors' liability to Customer, whether in contract, tort (including negligence), breach of warranty, or otherwise, exceed the price paid by Customer for the Software that gave rise to the claim, or if the Software is embedded in another Juniper product, the price paid by Customer for such other product. Customer acknowledges and agrees that Juniper has set its prices and entered into this Agreement in reliance upon the disclaimers of warranty and the limitations of liability set forth herein, that the same reflect an allocation of risk between the Parties (including the risk that a contract remedy may fail of its essential purpose and cause consequential loss), and that the same form an essential basis of the bargain between the Parties.

9. **Termination.** Any breach of this Agreement or failure by Customer to pay any applicable fees due shall result in automatic termination of the license granted herein. Upon such termination, Customer shall destroy or return to Juniper all copies of the Software and related documentation in Customer's possession or control.

10. **Taxes.** All license fees payable under this agreement are exclusive of tax. Customer shall be responsible for paying Taxes arising from the purchase of the license, or importation or use of the Software. If applicable, valid exemption documentation for each taxing jurisdiction shall be provided to Juniper prior to invoicing, and Customer shall promptly notify Juniper if their exemption is revoked or modified. All payments made by Customer shall be net of any applicable withholding tax. Customer will provide reasonable assistance to Juniper in connection with such withholding taxes by promptly: providing Juniper with valid tax receipts and other required documentation showing Customer's payment of any withholding taxes; completing appropriate applications that would reduce the amount of withholding tax to be paid; and notifying and assisting Juniper in any audit or tax proceeding related to transactions hereunder. Customer shall comply with all applicable tax laws and regulations, and Customer will promptly pay or reimburse Juniper for all costs and damages related to any liability incurred by Juniper as a result of Customer's non-compliance or delay with its responsibilities herein. Customer's obligations under this Section shall survive termination or expiration of this Agreement.

11. **Export.** Customer agrees to comply with all applicable export laws and restrictions and regulations of any United States and any applicable foreign agency or authority, and not to export or re-export the Software or any direct product thereof in violation of any such

restrictions, laws or regulations, or without all necessary approvals. Customer shall be liable for any such violations. The version of the Software supplied to Customer may contain encryption or other capabilities restricting Customer's ability to export the Software without an export license.

12. **Commercial Computer Software.** The Software is "commercial computer software" and is provided with restricted rights. Use, duplication, or disclosure by the United States government is subject to restrictions set forth in this Agreement and as provided in DFARS 227.7201 through 227.7202-4, FAR 12.212, FAR 27.405(b)(2), FAR 52.227-19, or FAR 52.227-14(ALT III) as applicable.

13. **Interface Information.** To the extent required by applicable law, and at Customer's written request, Juniper shall provide Customer with the interface information needed to achieve interoperability between the Software and another independently created program, on payment of applicable fee, if any. Customer shall observe strict obligations of confidentiality with respect to such information and shall use such information in compliance with any applicable terms and conditions upon which Juniper makes such information available.

14. **Third Party Software.** Any licensor of Juniper whose software is embedded in the Software and any supplier of Juniper whose products or technology are embedded in (or services are accessed by) the Software shall be a third party beneficiary with respect to this Agreement, and such licensor or vendor shall have the right to enforce this Agreement in its own name as if it were Juniper. In addition, certain third party software may be provided with the Software and is subject to the accompanying license(s), if any, of its respective owner(s). To the extent portions of the Software are distributed under and subject to open source licenses obligating Juniper to make the source code for such portions publicly available (such as the GNU General Public License ("GPL") or the GNU Library General Public License ("LGPL")), Juniper will make such source code portions (including Juniper modifications, as appropriate) available upon request for a period of up to three years from the date of distribution. Such request can be made in writing to Juniper Networks, Inc., 1194 N. Mathilda Ave., Sunnyvale, CA 94089, ATTN: General Counsel. You may obtain a copy of the GPL at http://www.gnu.org/licenses/gpl.html, and a copy of the LGPL at http://www.gnu.org/licenses/lgpl.html .

15. **Miscellaneous.** This Agreement shall be governed by the laws of the State of California without reference to its conflicts of laws principles. The provisions of the U.N. Convention for the International Sale of Goods shall not apply to this Agreement. For any disputes arising under this Agreement, the Parties hereby consent to the personal and exclusive jurisdiction of, and venue in, the state and federal courts within Santa Clara County, California. This Agreement constitutes the entire and sole agreement between Juniper and the Customer with respect to the Software, and supersedes all prior and contemporaneous agreements relating to the Software, whether oral or written (including any inconsistent terms contained in a purchase order), except that the terms of a separate written agreement executed by an authorized Juniper representative and Customer shall govern to the extent such terms are inconsistent or conflict with terms contained herein. No modification to this Agreement nor any waiver of any rights hereunder shall be effective unless expressly assented to in writing by the party to be charged. If any portion of this Agreement is held invalid, the Parties agree that such invalidity shall not affect the validity of the remainder of this Agreement. This Agreement and associated documentation has been written in the English language, and the Parties agree that the English version will govern. (For Canada: Les parties aux présentés confirment leur volonté que cette convention de même que tous les documents y compris tout avis qui s'y rattaché, soient redigés en langue anglaise. (Translation: The parties confirm that this Agreement and all related documentation is and will be in the English language)).

Table of Contents

Part 21 Class of Service

About This Guide

How to Use This Guide

This guide, the *Dell PowerConnect J-Series Ethernet Switch Complete Software Guide for Junos OS*, provides the following information about Juniper Networks Junos Operating System (Junos OS) for J-EX Series switches: software feature descriptions, configuration examples and tasks, management and monitoring instructions, and reference information.

To download the Dell PowerConnect J-EX Series documentation listed in Table 1 on page lxix, see the following Dell support website:

http://www.support.dell.com/manuals

Table 1: List of J-EX Series Guides for Junos OS Release 10.3

Title	Description
Dell PowerConnect J-Series J-EX4200 Ethernet Switch Hardware Guide	Component descriptions, site preparation, installation, replacement, and safety and compliance information for J-EX4200 switches
Dell PowerConnect J-Series J-EX8208 Ethernet Switch Hardware Guide	Component descriptions, site preparation, installation, replacement, and safety and compliance information for J-EX8208 switches
Dell PowerConnect J-Series J-EX8216 Ethernet Switch Hardware Guide	Component descriptions, site preparation, installation, replacement, and safety and compliance information for J-EX8216 switches
Dell PowerConnect J-Series Ethernet Switch Complete Software Guide for Junos OS	Software feature descriptions, configuration examples, and tasks for Juniper Networks Junos OS for J-EX Series switches

To download additional Junos OS documentation for J-EX Series and all other PowerConnect J-Series products, see the following Juniper Networks support website: http://www.juniper.net/techpubs/.

If the information in the latest release notes differs from the information in the documentation, follow the release notes.

Downloading Software

You can download Junos OS for J-EX Series switches from the Download Software area at http://www.support.dell.com/ . To download the software, you must have a Juniper Networks user account. For information about obtaining an account, see http://www.support.dell.com.

Documentation Conventions

Table 2: Notice Icons

Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.

Table 3: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the **configure** command: user@host> **configure**
`Fixed-width text like this`	Represents output that appears on the terminal screen.	`user@host>` **show chassis alarms** `No alarms currently active`
Italic text like this	• Introduces important new terms. • Identifies book names. • Identifies RFC and Internet draft titles.	• A policy *term* is a named structure that defines match conditions and actions. • *Junos OS System Basics Configuration Guide* • RFC 1997, *BGP Communities Attribute*
Italic text like this	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# **set system domain-name** *domain-name*
Plain text like this	Represents names of configuration statements, commands, files, and directories; IP addresses; configuration hierarchy levels; or labels on routing platform components.	• To configure a stub area, include the **stub** statement at the **[edit protocols ospf area area-id]** hierarchy level. • The console port is labeled **CONSOLE**.
< > (angle brackets)	Enclose optional keywords or variables.	**stub <default-metric** ***metric*****>;**

Table 3: Text and Syntax Conventions *(continued)*

Convention	Description	Examples
| (pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	**broadcast | multicast** (***string1*** **|** ***string2*** **|** ***string3***)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	**rsvp { # Required for dynamic MPLS only**
[ ] (square brackets)	Enclose a variable for which you can substitute one or more values.	**community name members [** ***community-ids*** **]**

Indention and braces ({ })	Identify a level in the configuration hierarchy.	[edit]
		routing-options {
		  static {
		  route default {
		  nexthop *address*;
		  retain;
		  }
		  }
		}

; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
J-Web GUI Conventions		
Bold text like this	Represents J-Web graphical user interface (GUI) items you click or select.	• In the Logical Interfaces box, select **All Interfaces**. • To cancel the configuration, click **Cancel**.
> (bold right angle bracket)	Separates levels in a hierarchy of J-Web selections.	In the configuration editor hierarchy, select **Protocols>Ospf**.

Repair and Warranty

CAUTION: Many repairs may only be done by a certified service technician. You should only perform troubleshooting and simple repairs as authorized in your product documentation, or as directed by the online or telephone service and support team. Damage due to servicing that is not authorized by Dell is not covered by your warranty. Read and follow the safety instructions that came with the product.

For more information, see the “Getting Help” chapter in your hardware guide.

Requesting Technical Support

For technical support, see http://www.support.dell.com. For more information, see “Getting Help” in the hardware guide for your Dell PowerConnect J-EX Series Ethernet Switch.

PART 1

Junos OS for J-EX Series Switches Product Overview

CHAPTER 1

Software Overview

J-EX Series Switch Software Features Overview

The following tables list the software features for J-EX Series Switches and the switches on which they are supported:

Table 4: Access Control Features

Feature	J-EX4200 Switches	J-EX8200 Switches
802.1X authentication	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Authentication fallback	Junos OS 10.3R1	Not supported
Captive portal authentication	Junos OS 10.2 or earlier	Not supported
Dynamic allocation of ternary content addressable memory (TCAM) memory to firewall filters	Junos OS 10.2 or earlier	Junos OS 10.3R1
Dynamic firewall filters for 802.1X authentication	Junos OS 10.2 or earlier	Not supported
Firewall filters and rate limiting	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
For a list of supported firewall filter match conditions and actions, see “Firewall Filter Match Conditions and Actions for J-EX Series Switches” on page 2728.		
Firewall filters on LAGs	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Firewall filter on loopback interface	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Firewall filters with IPv6	Junos OS 10.2 or earlier	Junos OS 10.3R1
MAC RADIUS authentication	Junos OS 10.2 or earlier	Junos OS 10.3R1
Policing	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Server fail fallback	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
TACACS+	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier

Table 5: Administration Features

Feature	J-EX4200 Switches	J-EX8200 Switches
System logging (syslog) over IPv6	Junos OS 10.2 or earlier	Not supported
System logging (syslog) over IPv4	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
System snapshot	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier

Table 6: Class-of-Service (CoS) Features

Feature	J-EX4200 Switches	J-EX8200 Switches
Class of service (CoS)—Class-based queuing with prioritization	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier

Table 6: Class-of-Service (CoS) Features *(continued)*

Feature	J-EX4200 Switches	J-EX8200 Switches
CoS—DSCP, IEEE 801.p, and IP precedence packet rewrites on routed VLAN interfaces (RVIs)	Junos OS 10.2 or earlier	Not supported
CoS—Interface-specific classifiers on routed VLAN interfaces (RVIs)	Junos OS 10.2 or earlier	Not supported
CoS—multidestination	Not applicable	Junos OS 10.2 or earlier
CoS support on LAGs	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
CoS support on routed VLAN interfaces (RVIs)	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Interface-specific CoS rewrite rules	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Junos OS EZQoS for CoS	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Port shaping and queue shaping	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Remarking of bridged packets	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier

Table 7: High Availability and Resiliency Features

Feature	J-EX4200 Switches	J-EX8200 Switches
Graceful protocol restart for BGP	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Graceful protocol restart for IS-IS	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Graceful protocol restart for OSPF	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Graceful Routing Engine switchover (GRES) for J-EX4200 Virtual Chassis configurations	Junos OS 10.2 or earlier	Not applicable
GRES for ARP entries	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
GRES for the forwarding database	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
GRES for port security	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Link Aggregation Control Protocol (LACP)	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
LACP support for dual-homing applications in data centers	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Link aggregation groups (LAGs)	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Link aggregation groups (LAGs) over Virtual Chassis ports (VCPs)	Junos OS 10.2 or earlier	Not applicable

Table 7: High Availability and Resiliency Features *(continued)*

Feature	J-EX4200 Switches	J-EX8200 Switches
Redundant trunk groups	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Virtual Router Redundancy Protocol (VRRP)	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
VRRP for IPv6 (except authentication type and authentication key)	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Virtual Chassis • Atomic software upgrade • Fast failover • Split and merge	Junos OS 10.2 or earlier	Not applicable
Virtual Chassis • Automatic software update on prospective member switches • Front-panel configuration of uplink module ports as Virtual Chassis ports (VCPs)	Junos OS 10.2 or earlier	Not applicable
Virtual Chassis • Autoprovisioning of Virtual Chassis ports (VCPs)	Junos OS 10.2 or earlier	Not applicable
Virtual Chassis • Support for SFP uplink module ports	Junos OS 10.2 or earlier	Not applicable

Table 8: Interfaces Features

Feature	J-EX4200 Switches	J-EX8200 Switches
Digital optical monitoring (DOM)	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Interface-range support	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Power over Ethernet (PoE)	Junos OS 10.2 or earlier	Not applicable
Power over Ethernet Plus (PoE+)	Not supported	Not supported
PoE power management mode	Junos OS 10.2 or earlier	Not supported
Unicast reverse-path forwarding (RPF)	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
VLAN-tagged Layer 3 subinterfaces	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier

Table 9: IP Address Management Features

Feature	J-EX4200 Switches	J-EX8200 Switches
DHCP server and relay with option 82 for Layer 2 VLANs	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
DHCPv6 and IPv6 DNS	Junos OS 10.2 or earlier	Not supported
Local DHCP server	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Static addresses	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier

Table 10: IPv6 Features

Feature	J-EX4200 Switches	J-EX8200 Switches
A separate software license is required for IPv6. See “Understanding Software Licenses for the J-EX Series Switch” on page 65.		
IPv6 (except multicast protocols)	Junos OS 10.2 or earlier	Not supported
IPv6 CoS (multi-field classification and rewrite, scheduling based on TC),	Junos OS 10.2 or earlier	Not supported
IPv6 multicast protocols (PIM, MLDv1/v2)	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
IPv6 Management and Services	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
IPv6 Path MTU Discovery	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier

Table 11: Layer 2 Network Protocols Features

Feature	J-EX4200 Switches	J-EX8200 Switches
802.1Q VLAN tagging	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
BPDU protection for spanning-tree protocols	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
GARP VLAN Registration Protocol (GVRP)	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Layer 2 protocol tunneling (L2PT)	Junos OS 10.2 or earlier	Not supported
Link Layer Discovery Protocol (LLDP)	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Link Layer Discovery Protocol-Media Endpoint Discovery (LLDP-MED) with voice over IP (VoIP) integration	Junos OS 10.2 or earlier	Not supported
Loop protection for spanning-tree protocols	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
MAC-based VLAN	Junos OS 10.2 or earlier	Not supported
Multiple VLAN Registration Protocol (MVRP)	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier

Table 11: Layer 2 Network Protocols Features *(continued)*

Feature	J-EX4200 Switches	J-EX8200 Switches
Private VLANs (PVLANs)	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Proxy ARP—restricted	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Proxy ARP—unrestricted	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Proxy ARP per VLAN	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Q-in-Q tunneling	Junos OS 10.2 or earlier	Not supported
Q-in-Q VLAN extended support for multiple S-VLANs per access interface, firewall-filter-based VLAN assignment, and routed VLAN interfaces (RVIs)	Junos OS 10.2 or earlier	Not supported
Root protection for spanning-tree protocols	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Spanning tree: • Spanning Tree Protocol (STP) • Rapid Spanning Tree Protocol (RSTP) • Multiple Spanning Tree Protocol (MSTP)	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Spanning tree: • VLAN Spanning Tree Protocol (VSTP)	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
RSTP and VSTP concurrent configuration	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Storm control	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Unknown Layer 2 unicast forwarding	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
VLAN ID translation	Junos OS 10.2 or earlier	Not supported
VLAN range	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier

Table 12: Layer 3 Protocols Features

Feature	J-EX4200 Switches	J-EX8200 Switches
Bidirectional Forwarding Detection (BFD)	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Border Gateway Protocol (BGP)	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
A separate software license is required for BGP and MBGP. See “Understanding Software Licenses for the J-EX Series Switch” on page 65.		
Filter-based forwarding	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier

Table 12: Layer 3 Protocols Features *(continued)*

Feature	J-EX4200 Switches	J-EX8200 Switches
Intermediate System-to-Intermediate System (IS-IS)	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
A separate software license is required for IS-IS. See "Understanding Software Licenses for the J-EX Series Switch" on page 65.		
IPv6 protocols: Open Shortest Path First version 3 (OSPFv3), RIPng, IS-IS for IPv6, IPv6 BGP	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
IPv6 Layer 3 multicast protocols	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Jumbo frames on routed VLAN interfaces (RVIs)	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Multicast Source Discovery Protocol (MSDP)	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
See the *Junos OS Routing Protocols Configuration Guide* at //www.juniper.net/techpubs/software/junos/index.html.		
OSPF Multitopology Routing (MT-OSPF)	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
See the *Junos OS Routing Protocols Configuration Guide* at //www.juniper.net/techpubs/software/junos/index.html.		
OSPFv2	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
OSPF IPsec support	Junos OS 10.3R1	Not supported
Routed VLAN interfaces (RVIs)	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Routing Information Protocol version 1 (RIPv1) and RIPv2	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Static routes	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Virtual routing and forwarding (VRF) with IPv4—virtual routing instances	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
VRF with IPv4—virtual routing instances for multicast traffic	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
VRF with IPv6—virtual routing instances for multicast traffic	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
VRF with IPv6—virtual routing instances for unicast traffic	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier

Table 13: MPLS Features

Feature	J-EX4200 Switches	J-EX8200 Switches
MPLS with RSVP-based label switched paths (LSPs) and MPLS-based circuit cross-connects (CCCs)	Junos OS 10.2 or earlier	Not supported
A separate software license is required for MPLS. See "Understanding Software Licenses for the J-EX Series Switch" on page 65.		

Table 13: MPLS Features *(continued)*

Feature	J-EX4200 Switches	J-EX8200 Switches
MPLS with class of service (CoS) and IP over MPLS	Junos OS 10.2 or earlier	Not supported

Table 14: Multicast Features

Feature	J-EX4200 Switches	J-EX8200 Switches
Internet Group Management Protocol (IGMP) version1 (v1) and IGMPv2	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
IGMPv3	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
IGMPv1/v2 snooping	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
IGMP snooping with routed VLAN interfaces (RVIs)	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
IGMPv3 snooping	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
IGMPv3 snooping EXCLUDE modes	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Multicast Service Discovery Protocol (MSDP)	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
See the *Junos OS Multicast Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.		
Multicast VLAN registration (MVR)	Junos OS 10.2 or earlier	Not supported
Protocol Independent Multicast dense mode (PIM DM)	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
See the *Junos OS Multicast Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.		
Protocol Independent Multicast source-specific multicast (PIM SSM)	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
See the *Junos OS Multicast Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.		
Protocol Independent Multicast sparse mode (PIM SM)	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
See the *Junos OS Multicast Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.		
Single-source multicast	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier

Table 15: Network Management and Monitoring Features

Feature	J-EX4200 Switches	J-EX8200 Switches
802.1ag Ethernet OAM connectivity fault management (CFM)	Junos OS 10.2 or earlier	Not supported

Table 15: Network Management and Monitoring Features *(continued)*

Feature	J-EX4200 Switches	J-EX8200 Switches
Ethernet OAM link fault management (LFM)	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Port mirroring	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Port mirroring enhancements • Layer 3 interface support • Multiple VLAN support	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Port mirroring enhancements • Ingress-only and egress-only attributes on VLAN members to avoid flooding mirrored traffic to member interfaces of a VLAN on the intermediate switch	Junos OS 10.2 or earlier	Not supported
Real-time performance monitoring (RPM)	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
RMON	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
sFlow monitoring technology	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
sFlow technology: Persistent IP addresses for agent IDs and use in datagrams	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Simple Network Management Protocol version 1 (SNMPv1), SNMPv2, and SNMPv3	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Time Domain Reflectometry (TDR)	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier

Table 16: Port Security Features

Feature	J-EX4200 Switches	J-EX8200 Switches
Automatic recovery for port error disable conditions	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
DHCP option 82	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
DHCP snooping	Junos OS 10.2 or earlier	Junos OS 10.3R1
Dynamic ARP inspection (DAI)	Junos OS 10.2 or earlier	Junos OS 10.3R1
IP source guard	Junos OS 10.2 or earlier	Junos OS 10.3R1
MAC limiting	Junos OS 10.2 or earlier	Junos OS 10.3R1
MAC move limiting	Junos OS 10.2 or earlier	Junos OS 10.3R1
Persistent storage for DHCP snooping	Junos OS 10.2 or earlier	Not supported

Table 16: Port Security Features *(continued)*

Feature	J-EX4200 Switches	J-EX8200 Switches
Static ARP support	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier

Table 17: System Management Features

Feature	J-EX4200 Switches	J-EX8200 Switches
Autoinstallation of configuration files	Junos OS 10.2 or earlier	Not supported
Automatic software download	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
Configuration rollback	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier
IP directed broadcast	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier

J-Web interface, for switch configuration and management
 	
NOTE:
To access the J-Web interface, your management device must have the following software installed:
• Operating system: Microsoft Windows XP Service Pack 3
• Browser version: One of the following. Other browsers might work but are not supported by J-Series platforms.
  ◦ Microsoft Internet Explorer version 7.0
  ◦ Mozilla Firefox version 3.0
• Additional requirements:
  ◦ Only English-language browsers are supported.
  ◦ The browser and the network must be able to receive and process HTTP/1.1 gzip compressed data.
	Junos OS 10.2 or earlier	Junos OS 10.2 or earlier

Online insertion and removal (OIR) of uplink modules	Junos OS 10.2 or earlier	Not applicable
Power budget management	Not applicable	Junos OS 10.2 or earlier

Related Documentation

- High Availability Features for J-EX Series Switches Overview on page 18
- Layer 3 Protocols Supported on J-EX Series Switches on page 13
- Layer 3 Protocols Not Supported on J-EX Series Switches on page 14
- J-EX4200 Switches Hardware Overview on page 25
- J-EX8208 Switch Hardware Overview on page 27
- J-EX8216 Switch Hardware Overview on page 30

Layer 3 Protocols Supported on J-EX Series Switches

J-EX Series switches support the Junos OS Layer 3 features and configuration statements listed in Table 18 on page 13:

Table 18: Supported Junos OS Layer 3 Protocol Statements and Features

Protocol	Notes	For More Information
BGP	Fully supported.	See the *Junos OS Routing Protocols Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.
BFD	Fully supported.	See the *Junos OS Routing Protocols Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.
ICMP	Fully supported.	See the *Junos OS Routing Protocols Configuration Guide* at http://www.juniper.net/techpubs/software/junos/
IGMPv1, v2 and v3	Fully supported.	See the *Junos OS Multicast Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.
IS-IS	Supported, with the exceptions noted in "Layer 3 Protocols Not Supported on J-EX Series Switches" on page 14.	See the *Junos OS Routing Protocols Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.
MLD	Supported (MLD versions 1 and 2	See the *Junos OS Multicast Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.
MPLS	Supported, with the exceptions noted in "Layer 3 Protocols Not Supported on J-EX Series Switches" on page 14.	See the *Junos OS MPLS Applications Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.
OSPFv1, v2 and v3	Supported, with the exceptions noted in "Layer 3 Protocols Not Supported on J-EX Series Switches" on page 14.	See the *Junos OS Routing Protocols Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.
PIM	Fully supported	See the *Junos OS Multicast Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.
RIP	Fully supported.	See the *Junos OS Routing Protocols Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.
RIPng	Fully supported.	See the *Junos OS Routing Protocols Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.
SNMP	Fully supported.	See the *Junos OS Network Management Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.

Table 18: Supported Junos OS Layer 3 Protocol Statements and Features *(continued)*

Protocol	Notes	For More Information
VRRP	Fully supported with exception of IPv6 support of VRRP on routed VLAN interfaces (RVIs).	See "Understanding VRRP on J-EX Series Switches" on page 1425. See also the *Junos OS High Availability Guide* at http://www.juniper.net/techpubs/software/junos/.

Related Documentation

- Layer 3 Protocols Not Supported on J-EX Series Switches on page 14
- J-EX Series Switch Software Features Overview on page 3

Layer 3 Protocols Not Supported on J-EX Series Switches

J-EX Series switches do not support the Junos OS Layer 3 protocols and features listed in Table 19 on page 14.

Table 19: Junos OS Layer 3 Protocol Statements and Features That Are Not Supported

Feature	Configuration Statements Not Supported on J-EX Series Switches
DVMRP	• **dvmrp** and subordinate statements
Flow aggregation (cflowd)	• **cflow** and subordinate statements
GRE	• Not supported
IPsec	• **[edit services]** statements related to IPsec
IS-IS: • ES-IS • IPv6 in multicast routing protocols	• **clns-routing** statement • **ipv6-multicast** statement • **lsp-interval** statement • **label-switched-path** statement • **lsp-lifetime** statement • **te-metric** statement
Logical routers	• **logical-routers** and subordinate statements

Table 19: Junos OS Layer 3 Protocol Statements and Features That Are Not Supported *(continued)*

Feature	Configuration Statements Not Supported on J-EX Series Switches
MPLS: • Fast Reroute (FRR) • Label Distribution Protocol (LDP) • Layer 3 VPNs • Multiprotocol BGP (MP-BGP) for VPN-IPv4 family • Pseudowire emulation (PWE3) • Routing policy statements related to Layer 3 VPNs and MPLS • Virtual Private LAN Service (VPLS)	• **ldp** and all subordinate statements
Network Address Translation (NAT)	• **nat** and subordinate statements • Policy statements related to NAT
OSPF	• **demand-circuit** statement • **label-switched-path** and subordinate statements • **neighbor** statement within an OSPF area • **peer-interface** and subordinate statements within an OSPF area • **sham-link** statement • **te-metric** statement
Routing instances: • Routing instance forwarding	• **l2vpn** and subordinate statements • **ldp** and subordinate statements • **vpls** and subordinate statements
SAP and SDP	• **sap** and all subordinate statements
General routing options in the **routing-options** hierarchy: • MPLS and label-switched-paths	• **auto-export** and subordinate statements • **dynamic-tunnels** and subordinate statements • **lsp-next-hop** and subordinate statements • **multicast** and subordinate statements • **p2mp-lsp-next-hop** and subordinate statements • **route-distinguisher-id** statement

Table 19: Junos OS Layer 3 Protocol Statements and Features That Are Not Supported *(continued)*

Feature	Configuration Statements Not Supported on J-EX Series Switches

Traffic sampling and fowarding in the **forwarding-options** hierarchy	• **accounting** and subordinate statements
	• **family mpls** and **family multiservice** under **hash-key** hierarchy
	• Under **monitoring** ***group-name*** **family inet output** hierarchy:
	  ◦ **cflowd** statement
	  ◦ **export-format-cflowd-version-5** statement
	  ◦ **flow-active-timeout** statement
	  ◦ **flow-export-destination** statement
	  ◦ **flow-inactive-timeout** statement
	  ◦ **interface** statement
	• **port-mirroring** statement (On J-EX Series switches, port mirroring is implemented using the **analyzer** statement.)
	• **sampling** and subordinate statements

Related Documentation

- Layer 3 Protocols Supported on J-EX Series Switches on page 13
- J-EX Series Switch Software Features Overview on page 3

Security Features for J-EX Series Switches Overview

Juniper Networks Junos operating system (Junos OS) is a network operating system that has been hardened through the separation of control forwarding and services planes, with each function running in protected memory. The control-plane CPU is protected by rate limiting, routing policy, and firewall filters to ensure switch uptime even under severe attack. In addition, the switches fully integrate with the Juniper Networks Unified Access Control (UAC) product to provide both standards-based 802.1X port-level access and Layer 2 through Layer 4 policy enforcement based on user identity. Access port security features such as dynamic Address Resolution Protocol (ARP) inspection, DHCP snooping, and MAC limiting are controlled through a single Junos OS CLI command.

J-EX Series Switches provide the following hardware and software security features:

Console Port—Allows use of the console port to connect to the Routing Engine through an RJ-45 cable. You then use the command-line interface (CLI) to configure the switch.

Out-of-Band Management—A dedicated management Ethernet port on the rear panel allows out-of-band management.

Software Images—All Junos OS images are signed by Juniper Networks certificate authority (CA) with public key infrastructure (PKI).

User Authentication, Authorization, and Accounting (AAA)—Features include:

- User and group accounts with password encryption and authentication.
- Access privilege levels configurable for login classes and user templates.

- RADIUS authentication, TACACS+ authentication, or both, for authenticating users who attempt to access the switch.
- Auditing of configuration changes through system logging or RADIUS/TACACS+.

802.1X Authentication—Provides network access control. Supplicants (hosts) are authenticated when they initially connect to a LAN. Authenticating supplicants before they receive an IP address from a DHCP server prevents unauthorized supplicants from gaining access to the LAN. J-EX Series switches support Extensible Authentication Protocol (EAP) methods, including EAP-MD5, EAP-TLS, EAP-TTLS, and EAP-PEAP.

Port Security—Access port security features include:

- DHCP snooping—Filters and blocks ingress DHCP server messages on untrusted ports; builds and maintains an IP-address/MAC-address binding database (called the DHCP snooping database).
- Dynamic ARP inspection (DAI)—Prevents ARP spoofing attacks. ARP requests and replies are compared against entries in the DHCP snooping database, and filtering decisions are made based on the results of those comparisons.
- MAC limiting—Protects against flooding of the Ethernet switching table.
- MAC move limiting—Detects MAC movement and MAC spoofing on access ports.
- Trusted DHCP server—With a DHCP server on a trusted port, protects against rogue DHCP servers sending leases.
- IP source guard—Mitigates the effects of IP address spoofing attacks on the Ethernet LAN. The source IP address in the packet sent from an untrusted access interface is validated against the source MAC address in the DHCP snooping database. The packet is allowed for further processing if the source IP address to source MAC address binding is valid; if the binding is not valid, the packet is discarded.
- DHCP option 82—Also known as the DHCP relay agent information option. Helps protect the J-EX Series switch against attacks such as spoofing (forging) of IP addresses and MAC addresses and DHCP IP address starvation. Option 82 provides information about the network location of a DHCP client, and the DHCP server uses this information to implement IP addresses or other parameters for the client.
- Unrestricted proxy ARP—The switch responds to all ARP messages with its own MAC address. Hosts that are connected to the switch's interfaces cannot communicate directly with other hosts. Instead, all communications between hosts go through the switch.
- Restricted proxy ARP—The switch does not respond to an ARP request if the physical networks of the source and target of the ARP request are the same. It does not matter whether the destination host has the same IP address as the incoming interface or a different (remote) IP address. An ARP request for a broadcast address elicits no reply.

Device Security—Storm control permits the switch to monitor unknown unicast and broadcast traffic and drop packets, or shut down, or temporarily disable the interface when a specified traffic level is exceeded, thus preventing packets from proliferating and degrading the LAN. You can enable storm control on access interfaces or trunk interfaces.

Firewall Filters—Allow auditing of various types of security violations, including attempts to access the switch from unauthorized locations. Firewall filters can detect such attempts and create audit log entries when they occur. The filters can also restrict access by limiting traffic to source and destination MAC addresses, specific protocols, or, in combination with policers, to specified data rates to prevent denial of service (DoS) attacks.

Policers—Provide rate-limiting capability to control the amount of traffic that enters an interface, which acts to counter DoS attacks.

Encryption Standards—Supported standards include:

- 128-, 192-, and 256-bit Advanced Encryption Standard (AES)
- 56-bit Data Encryption Standard (DES) and 168-bit 3DES

Related Documentation

- 802.1X for J-EX Series Switches Overview on page 2253
- Firewall Filters for J-EX Series Switches Overview on page 2721
- Port Security for J-EX Series Switches Overview on page 2545
- Understanding Proxy ARP on J-EX Series Switches on page 1059
- Understanding Storm Control on J-EX Series Switches on page 2511
- Understanding the Use of Policers in Firewall Filters on page 2752

High Availability Features for J-EX Series Switches Overview

High availability refers to the hardware and software components that provide redundancy and reliability for packet-based communications. This topic covers the following high availability features of J-EX Series Switches:

VRRP

You can configure the Virtual Router Redundancy Protocol (VRRP) or VRRP for IPv6 on Gigabit Ethernet interfaces, 10-Gigabit Ethernet interfaces, and logical interfaces on J-EX Series switches. When VRRP is configured, the switches act as virtual routing platforms. VRRP enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts. The VRRP routing platforms share the IP address corresponding to the default route configured on the hosts. At any time, one of the VRRP routing platforms is the master (active) and the others are backups. If the master routing platform fails, one of the backup routing platforms becomes the new master, providing a virtual default routing platform

and enabling traffic on the LAN to be routed without relying on a single routing platform. Using VRRP, a backup J-EX Series switch can take over a failed default switch within a few seconds. This is done with minimum loss of VRRP traffic and without any interaction with the hosts.

For more information on VRRP, see "Understanding VRRP on J-EX Series Switches" on page 1425.

Graceful Protocol Restart

With standard implementations of routing protocols, any service interruption requires an affected switch to recalculate adjacencies with neighboring switches, restore routing table entries, and update other protocol-specific information. An unprotected restart of a switch can result in forwarding delays, route flapping, wait times stemming from protocol reconvergence, and even dropped packets. Graceful protocol restart allows a restarting switch and its neighbors to continue forwarding packets without disrupting network performance. Because neighboring switches assist in the restart (these neighbors are called helper switches), the restarting switch can quickly resume full operation without recalculating algorithms from scratch.

On J-EX Series switches, graceful protocol restart can be applied to aggregate and static routes and for routing protocols (BGP, IS-IS, OSPF, and RIP).

Graceful protocol restart works similarly for the different routing protocols. The main benefits of graceful protocol restart are uninterrupted packet forwarding and temporary suppression of all routing protocol updates. Graceful protocol restart thus allows a switch to pass through intermediate convergence states that are hidden from the rest of the network. Most graceful restart implementations define two types of switches—the restarting switch and the helper switch. The restarting switch requires rapid restoration of forwarding state information so that it can resume the forwarding of network traffic. The helper switch assists the restarting switch in this process. Individual graceful restart configuration statements typically apply to either the restarting switch or the helper switch.

Redundant Routing Engines

Two to ten J-EX4200 switches can be interconnected to create a Virtual Chassis configuration that operates as a single network entity. Every Virtual Chassis configuration has a master and a backup. The master acts as the master Routing Engine and the backup acts as the backup Routing Engine. The Routing Engine provides the following functionality:

- Runs various routing protocols
- Provides the forwarding table to the Packet Forwarding Engines (PFEs) in all the member switches of the Virtual Chassis configuration
- Runs other management and control processes for the entire Virtual Chassis configuration

The master Routing Engine, which is in the master of the Virtual Chassis configuration, runs Junos OS in the master role. It receives and transmits routing information, builds and maintains routing tables, communicates with interfaces and Packet Forwarding

Engine components of the member switches, and has full control over the Virtual Chassis configuration.

The backup Routing Engine, which is in the backup of the Virtual Chassis configuration, runs Junos OS in the backup role. It stays in sync with the master Routing Engine in terms of protocol states, forwarding tables, and so forth. If the master becomes unavailable, the backup Routing Engine takes over the functions that the master Routing Engine performs.

Graceful Routing Engine Switchover

You can configure graceful Routing Engine switchover (GRES) in a Virtual Chassis configuration, allowing the configuration to switch from the master Routing Engine in the master to the backup Routing Engine in the backup with minimal interruption to network communications. When you configure GRES, the backup Routing Engine automatically synchronizes with the master Routing Engine to preserve kernel state information and forwarding state. Any updates to the master Routing Engine are replicated to the backup Routing Engine as soon as they occur. If the kernel on the master Routing Engine stops operating, the master Routing Engine experiences a hardware failure, or the administrator initiates a manual switchover, mastership switches to the backup Routing Engine.

When the backup Routing Engine assumes mastership in a redundant failover configuration (that is, when graceful Routing Engine switchover is not enabled), the Packet Forwarding Engines initialize their state to boot up state before they connect to the new master Routing Engine. In contrast, in a graceful switchover configuration, the Packet Forwarding Engines do not reinitialize their state, but resynchronize their state with the new master Routing Engine. The interruption to the traffic is minimal.

Virtual Chassis Software Upgrade and Failover Features

J-EX4200 switches provide these features for increased resiliency in Virtual Chassis configurations:

- Virtual Chassis atomic software upgrade—When you upgrade software in a Virtual Chassis configuration, the upgrade will either succeed or fail on all member switches, preventing the situation in which only some of the Virtual Chassis member switches are upgraded.
- Virtual Chassis fast failover—A hardware-assisted failover mechanism that automatically reroutes traffic and reduces traffic loss in the event of a link failure.
- Virtual Chassis split and merge—If there is a disruption to the Virtual Chassis configuration due to member switches failing or being removed from the configuration, the Virtual Chassis configuration splits into two separate Virtual Chassis.

Link Aggregation

You can combine multiple physical Ethernet ports to form a logical point-to-point link, known as a *link aggregation group (LAG)* or *bundle*. A LAG provides more bandwidth than a single Ethernet link can provide. Additionally, link aggregation provides network

redundancy by load-balancing traffic across all available links. If one of the links should fail, the system automatically load-balances traffic across all remaining links.

You can select up to eight Ethernet interfaces and include them within a LAG. In a J-EX4200 Virtual Chassis configuration, the interfaces that form a LAG can be on different members of the Virtual Chassis. See "Understanding Virtual Chassis Configurations and Link Aggregation" on page 702.

Related Documentation

- For more information on high availability features, see the *Junos OS High Availability Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.
- Virtual Chassis Overview on page 691
- Understanding Virtual Chassis Components on page 694
- Understanding Virtual Chassis Configurations and Link Aggregation on page 702
- Understanding VRRP on J-EX Series Switches on page 1425

Understanding Software Infrastructure and Processes

Each switch runs Junos OS for J-EX Series Switches on its general-purpose processors. Junos OS includes processes for Internet Protocol (IP) routing and for managing interfaces, networks, and the chassis.

Junos OS runs on the Routing Engine. The Routing Engine kernel coordinates communication among the Junos OS processes and provides a link to the Packet Forwarding Engine.

With the J-Web interface and the command-line interface (CLI) to Junos OS, you configure switching features and routing protocols and set the properties of network interfaces on your switch. After activating a software configuration, use either the J-Web or CLI user interface to monitor the switch, manage operations, and diagnose protocol and network connectivity problems.

NOTE:

To access the J-Web interface, your management device must have the following software installed:

- Operating system: Microsoft Windows XP Service Pack 3
- Browser version: One of the following. Other browsers might work but are not supported by J-Series platforms.
  - Microsoft Internet Explorer version 7.0
  - Mozilla Firefox version 3.0
- Additional requirements:
  - Only English-language browsers are supported.
  - The browser and the network must be able to receive and process HTTP/1.1 gzip compressed data.

- Routing Engine and Packet Forwarding Engine on page 22
- Junos OS Processes on page 23

Routing Engine and Packet Forwarding Engine

A switch has two primary software processing components:

- Packet Forwarding Engine—Processes packets; applies filters, routing policies, and other features; and forwards packets to the next hop along the route to their final destination.
- Routing Engine—Provides three main functions:
  - Creates the packet forwarding switch fabric for the switch, providing route lookup, filtering, and switching on incoming data packets, then directing outbound packets to the appropriate interface for transmission to the network

- Maintains the routing tables used by the switch and controls the routing protocols that run on the switch.
- Provides control and monitoring functions for the switch, including controlling power and monitoring system status.

Junos OS Processes

Junos OS running on the Routing Engine and Packet Forwarding Engine consists of multiple processes that are responsible for individual functions.

The separation of functions provides operational stability, because each process accesses its own protected memory space. In addition, because each process is a separate software package, you can selectively upgrade all or part of Junos OS, for added flexibility.

Table 20 on page 23 describes the primary Junos OS processes.

Table 20: Junos OS Processes

Process	Name	Description
Chassis process	chassisd	Detects hardware on the system that is used to configure network interfaces. Monitors the physical status of hardware components and field-replaceable units (FRUs), detecting when environment sensors such as temperature sensors are triggered. Relays signals and interrupts—for example, when devices are taken offline, so that the system can close sessions and shut down gracefully.
Ethernet switching process	eswd	Handles Layer 2 switching functionality such as MAC address learning, Spanning Tree protocol and access port security. The process is also responsible for managing Ethernet switching interfaces, VLANs, and VLAN interfaces. Manages Ethernet switching interfaces, VLANs, and VLAN interfaces.
Forwarding process	pfem	Defines how routing protocols operate on the switch. The overall performance of the switch is largely determined by the effectiveness of the forwarding process.
Interface process	dcd	Configures and monitors network interfaces by defining physical characteristics such as link encapsulation, hold times, and keepalive timers.
Management process	mgd	Provides communication between the other processes and an interface to the configuration database. Populates the configuration database with configuration information and retrieves the information when queried by other processes to ensure that the system operates as configured. Interacts with the other processes when commands are issued through one of the user interfaces on the switch. If a process terminates or fails to start when called, the management process attempts to restart it a limited number of times to prevent thrashing and logs any failure information for further investigation.

Table 20: Junos OS Processes *(continued)*

Process	Name	Description
Routing protocol process	rpd	Defines how routing protocols such as RIP, OSPF, and BGP operate on the device, including selecting routes and maintaining forwarding tables.

Related Documentation

- For more information about processes, see the *Junos OS Network Operations Guide* at http://www.juniper.net/techpubs/software/junos/.
- For more information about basic system parameters, supported protocols, and software processes, see the *Junos OS System Basics Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.

CHAPTER 2

Supported Hardware

J-EX4200 Switches Hardware Overview

J-EX Series Switches provide scalable connectivity for the enterprise market, including branch offices, campus locations, and data centers. The switches run under Junos OS, which provides Layer 2 and Layer 3 switching, routing, and security services. The same Junos OS code base that runs on J-EX Series switches also runs on all Dell PowerConnect J-SRX Series Services Gateways.

J-EX4200 Switches

J-EX4200 switches provide connectivity for medium- and high-density environments and scalability for growing networks. These switches can be deployed wherever you need a high density of Gigabit Ethernet ports (24 to 480 ports) or redundancy. Typically, J-EX4200 switches are used in large branch offices, campus wiring closets, and data centers where they can be positioned as the top device in a rack to provide connectivity for all the devices in the rack.

You can connect individual J-EX4200 switches together to form one unit and manage the unit as a single chassis, called a *Virtual Chassis*. You can add more member switches to the Virtual Chassis as needed, up to a total of 10 members.

J-EX4200 switches are available in models with 24 or 48 ports and with 8 ports equipped for PoE. All models provide ports that have 10/100/1000Base-T Gigabit Ethernet connectors and optional 1-gigabit small form-factor pluggable (SFP) transceivers or 10-gigabit small form-factor pluggable (SFP+) transceivers for use with fiber connections.

Additionally, a 24-port model provides 100Base-FX/1000Base-X SFP ports. This model is typically used as a small distribution switch.

All J-EX4200 switches have dedicated 64-Gbps Virtual Chassis ports that allow you to connect the switches to each other. You can also use optional uplink module ports to connect members of a Virtual Chassis across multiple wiring closets.

To provide carrier-class reliability, J-EX4200 switches include:

- Dual redundant power supplies that are field-replaceable and hot-swappable. An optional additional connection to an external power source is also available.
- A field-replaceable fan tray with three fans. The switch remains operational if a single fan fails.
- Redundant Routing Engines in a Virtual Chassis configuration. This redundancy enables GRES (graceful Routing Engine switchover) and nonstop active routing.
- Junos OS with its modular design that enables failed system processes to gracefully restart.

Uplink Modules

Optional uplink modules are available for all J-EX4200 switches. Uplink modules provide four 1-gigabit small form-factor pluggable (SFP) transceivers or two 10-gigabit small form-factor pluggable (SFP+) transceivers. You can use SFP or SFP+ ports to connect an access switch to a distribution switch or to interconnect member switches of a Virtual Chassis across multiple wiring closets.

Power over Ethernet (PoE) Ports

PoE ports provide electrical current to devices through the network cables so that separate power cords for devices such as IP phones, wireless access points, and security cameras are unnecessary. J-EX4200 switches have partial (8-port) PoE capability.

Related Documentation

- J-EX4200 Switch Models on page 26
- Field-Replaceable Units in J-EX4200 Switches
- Site Preparation Checklist for J-EX4200 Switches

J-EX4200 Switch Models

The J-EX4200 switch is available with 24 or 48 ports and with partial Power over Ethernet (PoE) capability. Table 21 on page 26 lists the J-EX4200 switch models.

Table 21: J-EX4200 Switch Models

Model	Ports	Number of PoE-enabled Ports	Power Supply (Minimum)
J-EX4200-24T	24 Gigabit Ethernet	First 8 ports	320 W
J-EX4200-48T	48 Gigabit Ethernet	First 8 ports	320 W
J-EX4200-24F	24 small form-factor pluggable (SFP) transceivers	Not applicable	320 W

Related Documentation

- Front Panel of a J-EX4200 Switch
- Rear Panel of a J-EX4200 Switch
- J-EX4200 Switches Hardware Overview on page 25

J-EX8208 Switch Hardware Overview

Dell PowerConnect J-Series J-EX8208 Ethernet Switches provide high performance, scalable connectivity, and carrier-class reliability for high-density environments such as campus-aggregation and data-center networks. The J-EX8208 switch is a modular system that provides high availability and redundancy for all major hardware components, including Routing Engines, switch fabric, fan tray, and power supplies.

You can manage J-EX8208 switches using the same Junos OS interfaces that you use for other Junos OS platforms—the Junos OS command-line interface (CLI) and the J-Web graphical interface.

Software

The Dell PowerConnect J-EX Series Switches—PowerConnect J-EX4200 Ethernet Switches and PowerConnect J-EX8200 Ethernet Switches—run under Junos OS, which provides Layer 2 and Layer 3 switching, routing, and security services. The same Junos OS code base that runs on J-EX Series switches also runs on Dell PowerConnect J-SRX Series Services Gateways.

Chassis Physical Specifications

The J-EX8208 switch is 14 rack units (14 U) in size (1/3 rack); three J-EX8208 switches can fit in a standard 42 U rack. Each J-EX8208 switch is designed to optimize rack space and cabling. See Figure 1 on page 28.

Figure 1: J-EX8208 Switch

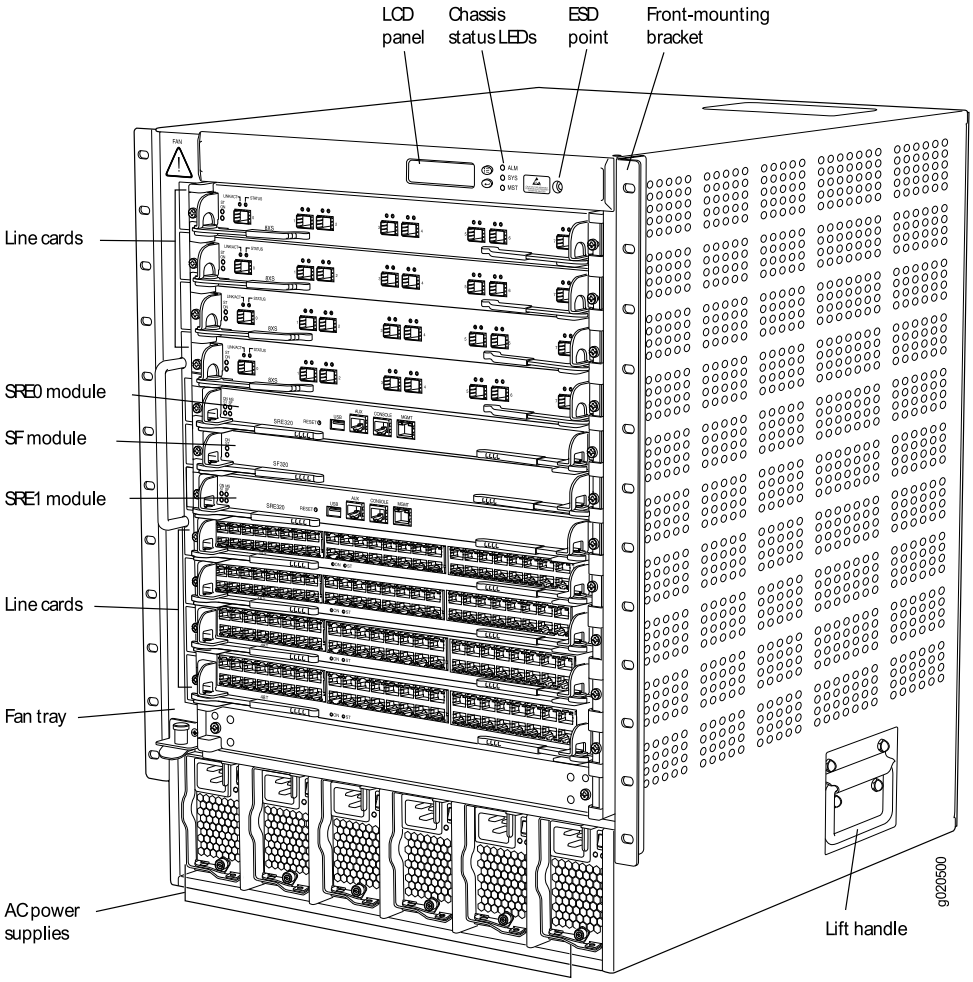

The J-EX8208 switch has a chassis-level LCD panel that displays Routing Engine and switch fabric status as well as chassis components' alarm information for rapid problem identification. The LCD panel provides a user-friendly interface for performing initial switch configuration, rolling back a configuration, or restoring the switch to its default settings. See LCD Panel in a J-EX8200 Switch.

The J-EX8208 chassis backplane distributes the data, control, and management signals to various system components along with distributing power throughout the system.

See Chassis Physical Specifications of a J-EX8208 Switch.

Routing Engines and Switch Fabric

Switching functionality, system management, and system control functions of a J-EX8208 switch are performed by the Switch Fabric and Routing Engine (SRE) module. See Switch Fabric and Routing Engine (SRE) Module in a J-EX8208 Switch. An SRE module contains a Routing Engine and switch fabric. The SRE modules are installed in the front of the chassis in the slots labeled SRE0 and SRE1. See Slot Numbering for a J-EX8208 Switch.

A base configuration J-EX8208 switch has one SRE module. A redundant configuration J-EX8208 switch has a second SRE module. See J-EX8208 Switch Configurations.

The Switch Fabric (SF) module, working with the SRE module, provides the necessary switching functionality to a base configuration J-EX8208 switch. The SF module is installed in the front of the chassis in the slot labeled SF. In a redundant configuration the SF module provides a redundant switch fabric. The additional switch fabric provides full 2+1 switch fabric redundancy to the switch. See Switch Fabric (SF) Module in a J-EX8208 Switch.

Line Cards

The J-EX8208 switch features eight horizontal line card slots and supports the line rate for each line card. The line cards in J-EX8200 switches combine a Packet Forwarding Engine and Ethernet interfaces on a single assembly. They are field-replaceable units (FRUs) that can be installed in the line card slots labeled 0 through 7 on the front of the switch chassis. See Slot Numbering for a J-EX8208 Switch. All line cards are hot-removable and hot-insertable.

The following line cards are available for J-EX8208 switches:

- 8-port 10-Gigabit Ethernet SFP+ line card: This line card has eight 10-gigabit SFP+ ports on its faceplate in which you can install SFP+ transceivers. See 8-port SFP+ Line Card in a J-EX8200 Switch.
- 48-port 100/1000 SFP line card: This line card has 48 1-gigabit SFP ports on its faceplate in which you can install SFP transceivers. See 48-port SFP Line Card in a J-EX8200 Switch.
- 48-port 10/100/1000 RJ-45 line card: This line card had 48 10/100/1000 Gigabit Ethernet ports with RJ-45 connectors on its faceplate. See 48-port RJ-45 Line Card in a J-EX8200 Switch.

Cooling System

The cooling system in a J-EX8208 switch consists of a hot-removable and hot-insertable fan tray. The fan tray contains 12 fans. The fan tray installs vertically on the left front of the chassis and provides side-to-side chassis cooling. See Cooling System and Airflow in a J-EX8208 Switch.

Power Supplies

Power supplies for the J-EX8208 switch are fully redundant, load-sharing, and hot-removable and hot-insertable field-replaceable units (FRUs). Each J-EX8208 switch chassis can hold up to six 2000 W AC power supplies.

The 2000 W AC power supplies support both low-voltage line (100–120 VAC) and high-voltage line (200–240 VAC) AC power configurations on a J-EX8208 switch. Each 2000 W AC power supply delivers 2000 W of power at high voltage (200–240 VAC) or 1200 W at low voltage (100–120 VAC) to the J-EX8208 chassis.

Only two AC power supplies are required for the base AC configuration and switch powerup. The redundant AC configuration ships with six AC power supplies to provide the capacity to power the system using N+1 or N+N power redundancy.

Related Documentation

- Field-Replaceable Units in a J-EX8208 Switch
- Connecting and Configuring a J-EX Series Switch (CLI Procedure) on page 161
- Connecting and Configuring a J-EX Series Switch (J-Web Procedure) on page 163

J-EX8216 Switch Hardware Overview

The Dell PowerConnect J-Series J-EX8216 Ethernet Switch is a half-rack, midplane architecture, modular Ethernet switch that is designed for ultra high-density environments such as campus aggregation, data center, or high performance core switching environments. J-EX8216 switches provide high-availability and redundancy for all major hardware components, including Routing Engine (RE) modules, Switch Fabric (SF) modules, fan trays (with redundant fans), and load-sharing 2000 W AC and 3000 W AC power supplies. Like other J-EX8200 Ethernet Switches, J-EX8216 switches provide high performance, scalable connectivity, and carrier-class reliability.

You can manage J-EX8216 switches using the same Junos OS interfaces that you use for other Junos OS platforms—the Junos OS command-line interface (CLI) and the J-Web graphical interface.

Software

The Dell PowerConnect J-EX Series Switches—PowerConnect J-EX4200 Ethernet Switches and PowerConnect J-EX8200 Ethernet Switches—run under Junos OS, which provides Layer 2 and Layer 3 switching, routing, and security services. The same Junos OS code base that runs on J-EX Series switches also runs on Dell PowerConnect J-SRX Series Services Gateways.

Chassis Physical Specifications, LCD Panel, and Midplane

J-EX8216 switches are designed to optimize rack space and cabling. The J-EX8216 switch is 21 rack units (21 U) in size (1/2 rack); two J-EX8216 switches can fit in a standard 42 U rack. See Figure 2 on page 31 and Figure 3 on page 32 and Chassis Physical Specifications of a J-EX8216 Switch.

Figure 2: J-EX8216 Switch Front

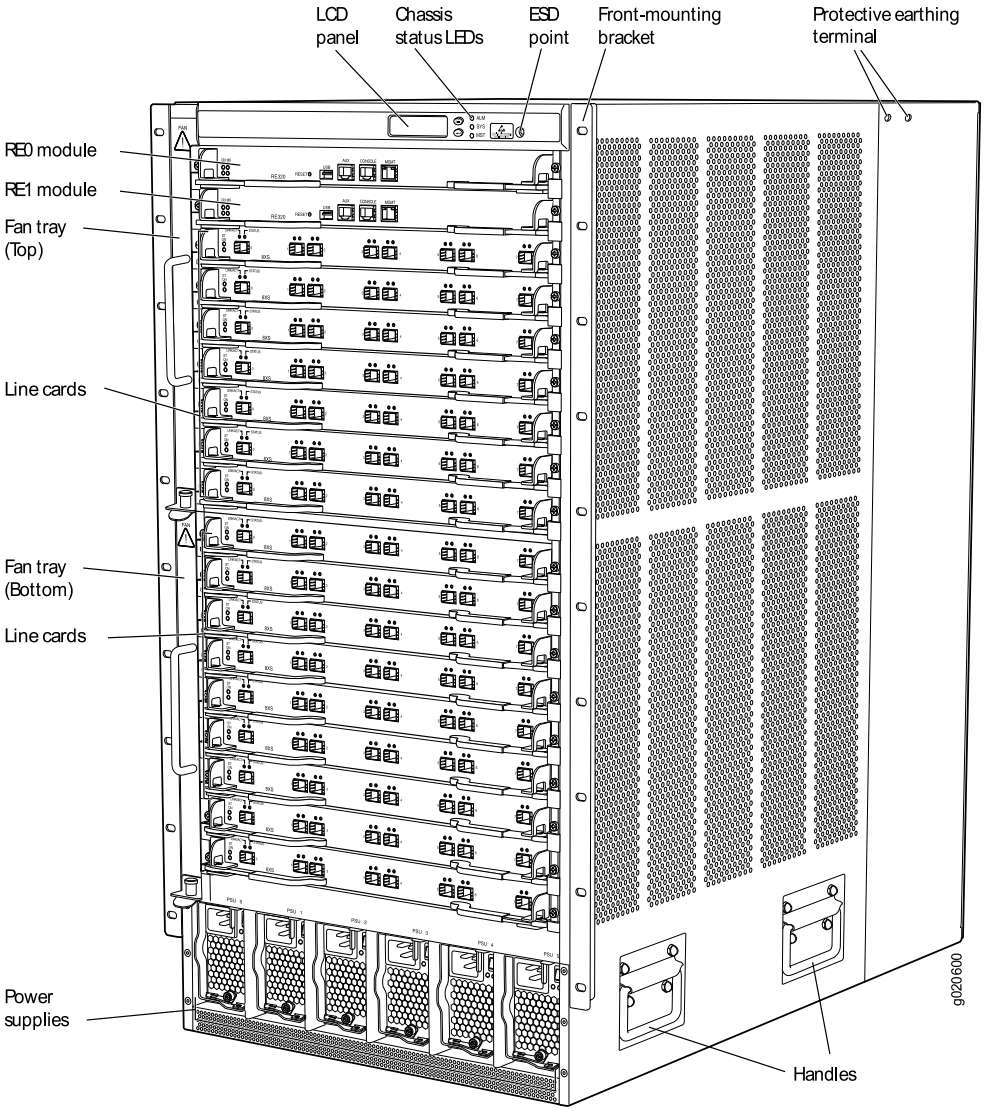

Figure 3: J-EX8216 Switch Rear

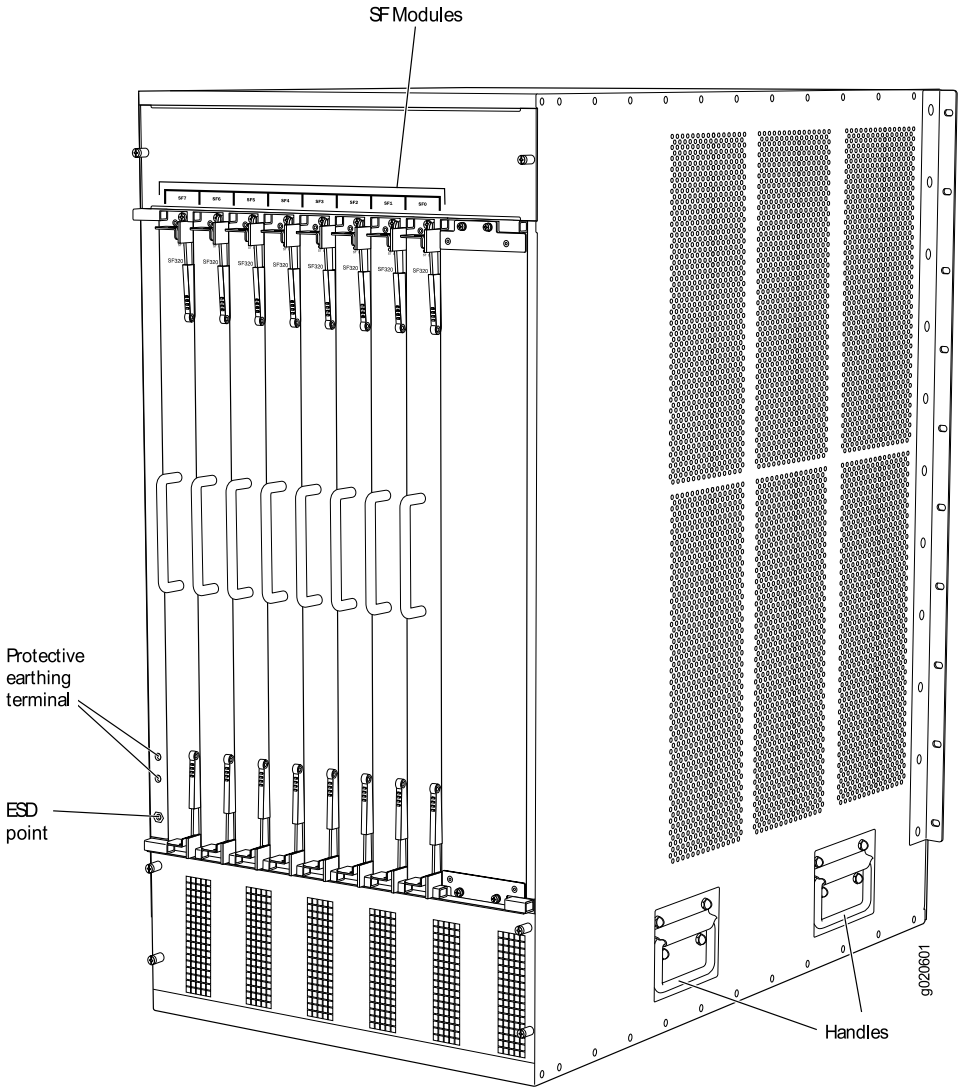

The J-EX8216 switch has a chassis-level LCD panel that displays Routing Engine and switch fabric status as well as chassis components' alarm information for rapid problem identification. The LCD panel provides a user-friendly interface for performing initial switch configuration, rolling back a configuration, or restoring the switch to the factory default configuration. See LCD Panel in a J-EX8200 Switch.

The J-EX8216 chassis midplane distributes the data, control, and management signals to system components and distributes power throughout the system. See Midplane in a J-EX8216 Switch.

Routing Engines and Switch Fabric

System management and system control functions of a J-EX8216 switch are performed by the Routing Engine (RE) module. An RE module contains a Routing Engine. The RE modules are hot-insertable and hot-removable field-replaceable units (FRUs) that are

installed in the front of the chassis in the slots labeled RE0 and RE1. A base configuration (AC version) J-EX8216 switch has one RE module. A redundant configuration J-EX8216 switch has a second RE module for redundancy. See Routing Engine (RE) Module in a J-EX8216 Switch and J-EX8216 Switch Configurations.

The Switch Fabric (SF) modules provide the switching functionality to a J-EX8216 switch. The SF modules are hot-insertable and hot-removable field-replaceable units (FRUs). All eight SF modules are installed in the rear of the chassis in the slots labeled SF7 through SF0. In a J-EX8216 switch, all eight SF modules are active and must be installed in the switch for normal operation. If a single SF module fails, the input/output traffic for that module is load-balanced among the remaining SF modules, providing graceful degradation in midplane performance. The impact of an SF module failure on the performance of a J-EX8216 switch varies based on the type of line cards installed in the switch and the traffic mix flowing through them. In a J-EX8216 switch configuration that is fully loaded with 8-port 10-Gigabit Ethernet SFP+ line cards, if one SF module fails, the remaining seven SF modules still have sufficient switching capacity to maintain continuous switch operation at full wire-rate performance. See Switch Fabric (SF) Modules in a J-EX8216 Switch.

Line Cards

The J-EX8216 switch features 16 horizontal line card slots and supports wire-rate performance for all packet sizes for the installed line cards. The line cards in J-EX8200 switches combine a Packet Forwarding Engine and Ethernet interfaces on a single assembly. They are field-replaceable units (FRUs), and you can install them in the slots labeled 0 through 15 on the front of the switch chassis. All line cards are hot-insertable and hot-removable.

The following line cards are available for J-EX8216 switches:

- 8-port 10-Gigabit Ethernet SFP+ line card: This line card has eight 10-gigabit SFP+ ports on its faceplate in which you can install SFP+ transceivers. See 8-port SFP+ Line Card in a J-EX8200 Switch.
- 48-port 100/1000 SFP line card: This line card has 48 1-gigabit SFP ports on its faceplate in which you can install SFP transceivers. See 48-port SFP Line Card in a J-EX8200 Switch.
- 48-port 10/100/1000 RJ-45 line card: This line card has 48 10/100/1000 Gigabit Ethernet ports with RJ-45 connectors on its faceplate. See 48-port RJ-45 Line Card in a J-EX8200 Switch.

Cooling System

The cooling system in a J-EX8216 switch consists of two hot-insertable and hot-removable, field-replaceable unit (FRU) fan trays. Each fan tray contains nine fans. Both fan trays install vertically on the left front of the chassis and provide side-to-side chassis cooling and front-to-side cooling. The top and bottom fan trays are identical and interchangeable. However, only the top fan tray cools the SF modules installed in the rear of the chassis. See Cooling System and Airflow in a J-EX8216 Switch.

Power Supplies

Power supplies for the J-EX8216 switch are fully redundant, load-sharing, and hot-insertable and hot-removable field-replaceable units (FRUs). Each J-EX8216 switch chassis can hold up to six 2000 W AC or six 3000 W AC power supplies.

The 2000 W AC power supplies support both low-voltage line (100–120 VAC) and high-voltage line (200–240 VAC) AC power configurations on a J-EX8216 switch.

Each 3000 W AC power supply delivers 3000 W of power at high voltage (200–240 VAC) to the J-EX8216 chassis. Low-voltage input is not supported for the 3000 W AC power supplies on the J-EX8216 switch.

The redundant AC configuration ships with six AC power supplies to provide the capacity to power the system using N+1 or N+N power redundancy. See AC Power Supply in a J-EX8200 Switch and J-EX8216 Switch Configurations.

CAUTION: Mixing different types of power supplies in the same chassis is not a supported configuration.

Related Documentation

- Field-Replaceable Units in a J-EX8216 Switch
- Slot Numbering for a J-EX8216 Switch
- Connecting and Configuring a J-EX Series Switch (CLI Procedure) on page 161
- Connecting and Configuring a J-EX Series Switch (J-Web Procedure) on page 163

PART 2

Complete Software Configuration Statement Hierarchy

CHAPTER 3

Complete Software Configuration Statement Hierarchy

[edit access] Configuration Statement Hierarchy

```
access {
  profile profile-name {
    accounting {
      order [ radius | none ];
      accounting-stop-on-access-deny;
      accounting-stop-on-failure;
    }
    authentication-order [ authentication-method ];
    radius {
      accounting-server [ server-address ];
      authentication-server [ server-address ];
    }
  }
}
```

Related Documentation

- Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch on page 2267

- Configuring 802.1X RADIUS Accounting (CLI Procedure) on page 2339

[edit chassis] Configuration Statement Hierarchy

```
chassis {
  aggregated-devices {
    ethernet {
      device-count number;
    }
  }
  auto-image-upgrade;
  }
  fpc slot {
    pic pic-number {
      sfpplus {
        pic-mode mode;
    power-budget-priority priority;
    }
  lcd-menu fpc slot-number {
    menu-item (menu-name | menu-option);
  }
}
  psu {
    redundancy {
      n-plus-n;
    }
  }
  redundancy{
    graceful-switchover ;
  }
```

Related Documentation

- Understanding Aggregated Ethernet Interfaces and LACP on page 867
- Understanding Power Management on J-EX Series Switches on page 302

[edit class-of-service] Configuration Statement Hierarchy

```
class-of-service {
  classifiers {
    (dscp | ieee-802.1 | inet-precedence) classifier-name {
      import (classifier-name | default);
      forwarding-class class-name {
        loss-priority loss-priority {
          code-points [ aliases ] [ 6 bit-patterns ];
        }
      }
    }
  }
  code-point-aliases {
    (dscp | ieee-802.1 | inet-precedence) {
      alias-name bits;
    }
  }
```

```
    forwarding-classes {
      class class-name queue-num queue-number priority ( high | low );
    }
    interfaces {
      interface-name  {
        scheduler-map map-name;
        unit logical-unit-number {
          forwarding-class class-name;
          classifiers {
            (dscp | ieee-802.1 | inet-precedence) (classifier-name | default);
          }
        }
      }
    }
    multi-destination {
      family {
        ethernet {
          broadcast forwarding-class-name;
        }
        inet {
          classifiers {
            (dscp |inet-precedence) classifier-name;
          }
        }
      }
      scheduler-map  map-name;
    }
    rewrite-rules {
      (dscp | ieee-802.1 | inet-precedence) rewrite-name {
        import (rewrite-name | default);
        forwarding-class class-name {
          loss-priority loss-priority code-point (alias | bits);
        }
      }
    }
    scheduler-maps {
      map-name {
        forwarding-class class-name scheduler scheduler-name;
      }
    }
    schedulers {
      scheduler-name  {
        buffer-size (percent percentage | remainder);
        drop-profile-map loss-priority loss-priority protocol protocol drop-profile
          profile-name;
        priority priority;
        shaping-rate (rate | percent percentage);
        transmit-rate (rate | percent percentage | remainder);
      }
    }
  }
```

Related Documentation

- Example: Configuring CoS on J-EX Series Switches on page 2883
- Defining CoS Code-Point Aliases (CLI Procedure) on page 2914 or Defining CoS Code-Point Aliases (J-Web Procedure) on page 2912

- Defining CoS Classifiers (CLI Procedure) on page 2914 or Defining CoS Classifiers (J-Web Procedure) on page 2916
- Defining CoS Forwarding Classes (CLI Procedure) on page 2918 or Defining CoS Forwarding Classes (J-Web Procedure) on page 2918
- Configuring CoS Tail Drop Profiles (CLI Procedure) on page 2925
- Defining CoS Schedulers (CLI Procedure) on page 2920 or Defining CoS Schedulers (J-Web Procedure) on page 2920
- Defining CoS Rewrite Rules (CLI Procedure) on page 2925 or Defining CoS Rewrite Rules (J-Web Procedure) on page 2926
- Assigning CoS Components to Interfaces (CLI Procedure) on page 2928 or Assigning CoS Components to Interfaces (J-Web Procedure) on page 2928

[edit ethernet-switching-options] Configuration Statement Hierarchy

```
ethernet-switching-options {
  analyzer {
    name {
      loss-priority priority;
      ratio number;
      input {
        ingress {
          interface (all | interface-name);
          vlan (vlan-id | vlan-name);
        }
        egress {
          interface (all | interface-name);
        }
      output {
        interface interface-name;
        vlan (vlan-id | vlan-name);
      }
    }
  }
  bpdu-block {
    disable-timeout timeout;
    interface (all | [interface-name]);
  }
  dot1q-tunneling {
    ether-type (0x8100 | 0x88a8 | 0x9100);
  }
  interfaces interface-name {
    no-mac-learning;
  }
  mac-notification {
    notification-interval seconds;
  }
  mac-table-aging-time seconds;
  port-error-disable {
    disable-timeout timeout;
  }
```

```
redundant-trunk-group {
  group-name name {
    interface interface-name <primary>;
  }
}
secure-access-port {
  dhcp-snooping-file {
    location local_pathname | remote_URL;
    timeout seconds;
    write-interval seconds;
  }
  interface (all | interface-name) {
    allowed-mac {
      mac-address-list;
    }
    (dhcp-trusted | no-dhcp-trusted );
    mac-limit limit action action;
    no-allowed-mac-log;
    static-ip  ip-address {
      vlan  vlan-name;
      mac mac-address;
    }
  }
  vlan (all | vlan-name) {
    (arp-inspection | no-arp-inspection );
    dhcp-option82 {
      circuit-id {
        prefix hostname;
        use-interface-description;
        use-vlan-id;
      }
      remote-id {
        prefix hostname | mac | none;
        use-interface-description;
        use-string string;
      }
      vendor-id [string];
    }
    (examine-dhcp | no-examine-dhcp );
    (ip-source-guard | no-ip-source-guard);
    mac-move-limit limit action action;
  }
}
storm-control {
  action-shutdown;
  interface (all | interface-name) {
    bandwidth bandwidth;
    no-broadcast;
    no-unknown-unicast;
  }
}
traceoptions {
  file filename <files number> <no-stamp> <replace> <size size> <world-readable |
    no-world-readable>;
  flag flag <disable>;
}
```

```
    unknown-unicast-forwarding {
      vlan (all | vlan-name) {
        interface interface-name;
      }
    }
    voip {
      interface (all | [interface-name | access-ports]) {
        vlan vlan-name ;
        forwarding-class <assured-forwarding | best-effort | expedited-forwarding |
          network-control>;
      }
    }
}
```

Related Documentation

- Understanding Port Mirroring on J-EX Series Switches on page 3245
- Port Security for J-EX Series Switches Overview on page 2545
- Understanding BPDU Protection for STP, RSTP, and MSTP on J-EX Series Switches on page 1278
- Understanding Redundant Trunk Links on J-EX Series Switches on page 1049
- Understanding Storm Control on J-EX Series Switches on page 2511
- Understanding 802.1X and VoIP on J-EX Series Switches on page 2263
- Understanding Q-in-Q Tunneling on J-EX Series Switches on page 1051
- Understanding Unknown Unicast Forwarding on J-EX Series Switches on page 2512
- Understanding MAC Notification on J-EX Series Switches on page 1060

[edit firewall] Configuration Statement Hierarchy

```
firewall {
  family family-name {
    filter filter-name {
      interface-specific;
      term term-name {
        from {
          match-conditions;
        }
        then {
          action;
          action-modifiers;
        }
      }
    }
  }
  policer policer-name {
    filter-specific;
    if-exceeding {
      bandwidth-limit bps;
      burst-size-limit bytes;
    }
    then {
```

```
        policer-action;
      }
    }
  }
```

Related Documentation

- Firewall Filter Configuration Statements Supported by Junos OS for J-EX Series Switches on page 2806
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Configuring Firewall Filters (CLI Procedure) on page 2779
- Configuring Policers to Control Traffic Rates (CLI Procedure) on page 2788
- Firewall Filters for J-EX Series Switches Overview on page 2721

[edit forwarding-options] Configuration Statement Hierarchy

```
helpers {
  bootp {
    dhcp-option82 {
      circuit-id {
        prefix hostname;
        use-interface-description;
        use-vlan-id;
      }
      remote-id {
        prefix hostname | mac | none;
        use-interface-description;
        use-string string;
      }
      vendor-id <string>;
    }
    interface interface-name {
      dhcp-option82 {
        circuit-id {
          prefix hostname;
          use-interface-description;
          use-vlan-id;
        }
        remote-id {
          prefix hostname | mac | none;
          use-interface-description;
          use-string string;
        }
        vendor-id <string>;
      }
      source-address-giaddr;
    }
    source-address-giaddr;
  }
}
```

Related Documentation

- Example: Setting Up DHCP Option 82 with a J-EX Series Switch as Relay Agent Between Clients and a DHCP Server on page 2615
- Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 2646
- Understanding DHCP Option 82 for Port Security on J-EX Series Switches on page 2560
- DHCP/BOOTP Relay for J-EX Series Switches Overview on page 446
- For more information about the **[edit forwarding-options]** hierarchy and all its options, see the *Junos OS Policy Framework Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.

[edit interfaces] Configuration Statement Hierarchy

```
interfaces {
  aex {
    aggregated-ether-options {
      (flow-control | no-flow-control);
      lacp mode {
        periodic interval;
      }
      link-speed speed;
      minimum-links number;
    }
    description text;
    disable;
    hold-time up milliseconds down milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;
    traceoptions;
    (traps | no-traps);
    unit logical-unit-number {
      description text;
      disable;
      family family-name {...}
      proxy-arp (restricted | unrestricted);
      (traps | no-traps);
      vlan-id vlan-id-number;
    }
    vlan-tagging;
  }
  fe-fpc/pic/port {
    description text;
    disable;
    mtu bytes;
    no-gratuitous-arp-request;
    speed speed;
    traceoptions;
    (traps | no-traps);
    unit logical-unit-number {
      description text;
      disable;
      family family-name {...}
```

```
    proxy-arp (restricted | unrestricted);
    (traps | no-traps);
    vlan-id vlan-id-number;
  }
  vlan-tagging;
}
ge-fpc/pic/port {
  description text;
  disable;
  ether-options {
    802.3ad aex {
      lacp {
        force-up;
      }
    }
    (auto-negotiation | no-auto-negotiation);
    (flow-control | no-flow-control);
    link-mode mode;
    speed (auto-negotiation | speed);
  }
  hold-time up milliseconds down milliseconds;
  mtu bytes;
  no-gratuitous-arp-request;
  traceoptions;
  (traps | no-traps);
  unit logical-unit-number {
    description text;
    disable;
    family family-name {...}
    proxy-arp (restricted | unrestricted);
    rpm;
    (traps | no-traps);
    vlan-id vlan-id-number;
  }
  vlan-tagging;
}
interface-range  interface-range name {
  description text;
  disable;
  ether-options {
    802.3ad aex {
      lacp {
        force-up;
      }
    }
    (auto-negotiation | no-auto-negotiation);
    (flow-control| no-flow-control);
    link-mode mode;
    speed (auto-negotiation | speed);
  }
  hold-time up milliseconds down milliseconds;
  member interface-name;
  member-range  starting-interface name to  ending-interface name;
  mtu bytes;
  unit logical-unit-number {
    description text;
```

```
      disable;
      family family-name {...}
      proxy-arp (restricted | unrestricted);
      rpm;
      (traps | no-traps);
      vlan-id vlan-id-number;
    }
  }
  lo0 {
    description text;
    disable;
    hold-time up milliseconds down milliseconds;
    traceoptions;
    (traps | no-traps);
    unit logical-unit-number {
      description text;
      disable;
      family family-name {...}
      (traps | no-traps);
    }
  }
  me0 {
    description text;
    disable;
    hold-time up milliseconds down milliseconds;
    no-gratuitous-arp-request;
    traceoptions;
    (traps | no-traps);
    unit logical-unit-number {
      description text;
      disable;
      family family-name {...}
      (traps | no-traps);
      vlan-id vlan-id-number;
    }
    vlan-tagging;
  }
  vlan {
    description  text;
    disable;
    hold-time up milliseconds down milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;
    traceoptions;
    (traps | no-traps);
    unit logical-unit-number {
      description text;
      disable;
      family family-name {...}
      proxy-arp (restricted | unrestricted);
      (traps | no-traps);
    }
  }
  vme {
    description text;
    disable;
```

```
    hold-time up milliseconds down milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;
    traceoptions;
    (traps | no-traps);
    unit logical-unit-number {
      description text;
      disable;
      family family-name {...}
      (traps | no-traps);
      vlan-id vlan-id-number;
    }
    vlan-tagging;
  }
  xe-fpc/pic/port {
    description text;
    disable;
    ether-options {
      802.3ad aex {
        lacp (802.3ad) {
          force-up;
        }
      }
      (auto-negotiation | no-auto-negotiation);
      (flow-control | no-flow-control);
      link-mode mode;
      speed (auto-negotiation | speed);
    }
    hold-time up milliseconds down milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;
    traceoptions;
    (traps | no-traps);
    unit logical-unit-number {
      description text;
      disable;
      family family-name {...}
      proxy-arp (restricted | unrestricted);
      rpm;
      (traps | no-traps);
      vlan-id vlan-id-number;
    }
    vlan-tagging;
  }
}
```

Related Documentation

- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 919
- Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 922
- Configuring a Layer 3 Subinterface (CLI Procedure) on page 930
- Configuring Routed VLAN Interfaces (CLI Procedure) on page 1137
- Configuring the Virtual Management Ethernet Interface for Global Management of a Virtual Chassis (CLI Procedure) on page 797

- J-EX Series Switches Interfaces Overview on page 863
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

[edit poe] Configuration Statement Hierarchy

```
poe {
  guard-band watts;
  interface (all | interface-name) {
    disable;
    maximum-power watts;
    priority (high | low);
    telemetries {
      disable;
      duration hours;
      interval minutes;
    }
  }
  management (class | static);
  notification-control {
    fpc slot-number {
      disable;
    }
  }
}
```

Related Documentation

- Example: Configuring PoE Interfaces with Different Priorities on a J-EX Series Switch on page 3015
- Configuring PoE (CLI Procedure) on page 3021
- PoE and J-EX Series Switches Overview on page 3009

[edit protocols] Configuration Statement Hierarchy

```
protocols {
  connections {
    remote-interface-switch connection-name {
      interface interface-name.unit-number;
      transmit-lsp label-switched-path;
      receive-lsp label-switched-path;
    }
  }
  dot1x {
    authenticator {
      authentication-profile-name profile-name;
      interface (all | [ interface-names ]) {
        disable;
        guest-vlan ( vlan-id | vlan-name);
        mac-radius <restrict>;
        maximum-requests number;
        no-reauthentication;
```

```
        quiet-period seconds;
        reauthentication {
          interval seconds;
        }
        retries number;
        server-fail (deny | permit | use-cache | vlan-id | vlan-name);
        server-reject-vlan (vlan-id | vlan-name);
        server-timeout seconds;
        supplicant (multiple | single | single-secure);
        supplicant-timeout seconds;
        transmit-period seconds;
      }
      static mac-address {
        interface interface-name;
        vlan-assignment (vlan-id |vlan-name);
      }
  }
  gvrp {
    <enable | disable>;
    interface (all | [interface-name]) {
      disable;
    }
    join-timer millseconds;
    leave-timer milliseconds;
    leaveall-timer milliseconds;
  }
  igmp-snooping {
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>
        <match regex>;
      flag flag (detail | disable | receive | send);
    }
    vlan (vlan-id | vlan-number) {
      data-forwarding {
        source {
          groups group-prefix;
        }
        receiver {
          source-vlans vlan-list;
          install ;
        }
      }
      disable {
        interface interface-name
      }
      immediate-leave;
      interface interface-name {
        group-limit limit;
        multicast-router-interface;
        static {
          group ip-address;
        }
      }
      proxy ;
      query-interval seconds;
      query-last-member-interval seconds;
```

```
    query-response-interval seconds;
    robust-count number;
  }
}
lldp {
  disable;
  advertisement-interval seconds;
  hold-multiplier number;
  interface (all | interface-name) {
    disable;
  }
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>
      <match regex>;
    flag flag (detail | disable | receive | send);
  }
}
lldp-med {
  disable;
  fast-start number;
  interface (all | interface-name) {
    disable;
    location {
      elin number;
      civic-based {
        what number;
        country-code code;
        ca-type {
          number {
            ca-value value;
          }
        }
      }
    }
  }
}
mpls {
  interface ( all | interface-name );
  label-switched-path lsp-name to remote-provider-edge-switch;
  path destination {
    <address | hostname> <strict | loose>
  }
mstp {
  disable;
  bpdu-block-on-edge;
  bridge-priority priority;
  configuration-name name;
  forward-delay seconds;
  hello-time seconds;
  interface (all | interface-name) {
    disable;
    bpdu-timeout-action {
      block;
      alarm;
    }
    cost  cost;
```

```
    edge;
    mode mode;
    no-root-port;
    priority priority;
  }
  max-age seconds;
  max-hops hops;
  msti msti-id {
    vlan (vlan-id | vlan-name);
    interface interface-name {
      disable;
      cost  cost;
      edge;
      mode mode;
      priority priority;
    }
  }
  revision-level revision-level;
  traceoptions {
    file filename <files number > <size size> <no-stamp | world-readable |
      no-world-readable>;
    flag flag;
  }
}
mvrp {
  disable
  interface (all | interface-name) {
    disable;
    join-timer milliseconds;
    leave-timer milliseconds;
    leaveall-timer milliseconds;
    registration (forbidden | normal);
  }
  no-dynamic-vlan;
  traceoptions {
    file filename <files number > <size size> <no-stamp | world-readable |
      no-world-readable>;
    flag flag;
  }
}
oam {
  ethernet{
    connectivity-fault-management {
      action-profile profile-name {
        default-actions {
          interface-down;
        }
      }
      linktrace {
        age (30m | 10m | 1m | 30s | 10s);
        path-database-size path-database-size;
      }
      maintenance-domain domain-name {
        level number;
        mip-half-function (none | default |explicit);
        name-format (character-string | none | dns | mac+2oct);
```

```
            maintenance-association ma-name {
              continuity-check {
                hold-interval minutes;
                interval (10m | 10s | 1m | 1s| 100ms);
                loss-threshold number;
              }
              mep mep-id {
                auto-discovery;
                direction down;
                interface interface-name;
                remote-mep mep-id {
                  action-profile profile-name;
                }
              }
            }
          }
        }
        link-fault-management {
          action-profile profile-name;
            action {
              syslog;
              link-down;
            }
            event {
              link-adjacency-loss;
              link-event-rate;
              frame-error count;
              frame-period count;
              frame-period-summary count;
              symbol-period count;
            }
            interface interface-name {
              link-discovery (active | passive);
              pdu-interval interval;
              event-thresholds threshold-value;
              remote-loopback;
              event-thresholds {
                frame-errorcount;
                frame-period count;
                frame-period-summary count;
                symbol-period count;
              }
            }
            negotiation-options {
              allow-remote-loopback;
              no-allow-link-events;
            }
        }
      }
    }
    rstp {
      disable;
      bpdu-block-on-edge;
      bridge-priority priority;
      forward-delay seconds;
      hello-time seconds;
```

```
    interface (all | interface-name) {
      disable;
      bpdu-timeout-action {
        block;
        alarm;
      }
      cost  cost;
      edge;
      mode mode;
      no-root-port;
      priority priority;
    }
    max-age seconds;
    }
    traceoptions {
      file filename <files number > <size size> <no-stamp | world-readable |
        no-world-readable>;
      flag flag;
    }
  }
  sflow {
    agent-id
    collector {
      ip-address;
      udp-port port-number;
    }
    disable;
    interfaces interface-name {
      disable;
      polling-interval seconds;
      sample-rate number;
    }
    polling-interval seconds;
    sample-rate number;
    source-ip
  }
  stp {
    disable;
    bridge-priority priority;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
      disable;
      bpdu-timeout-action {
        block;
        alarm;
      }
      cost  cost;
      edge;
      mode mode;
      no-root-port;
      priority priority;
    }
    max-age seconds;
  }
    traceoptions {
```

```
        file filename <files number > <size size> <no-stamp | world-readable |
          no-world-readable>;
        flag flag;
      }
    vstp {
      bpdu-block-on-edge;
      disable;
      force-version stp;
      vlan (all | vlan-id | vlan-name) {
        bridge-priority priority;
        forward-delay seconds;
        hello-time seconds;
        interface (all | interface-name) {
          bpdu-timeout-action {
            alarm;
            block;
          }
          cost  cost;
          disable;
          edge;
          mode mode;
          no-root-port;
          priority priority;
        }
        max-age seconds;
        traceoptions {
          file filename <files number > <size size> <no-stamp | world-readable |
            no-world-readable>;
          flag flag;
        }
      }
    }
  }
```

Related Documentation

- 802.1X for J-EX Series Switches Overview on page 2253
- Example: Configure Automatic VLAN Administration Using GVRP on page 1087
- Understanding MAC RADIUS Authentication on J-EX Series Switches
- Understanding Server Fail Fallback and 802.1X Authentication on J-EX Series Switches on page 2258
- IGMP Snooping on J-EX Series Switches Overview on page 2047
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261
- Understanding MSTP for J-EX Series Switches on page 1277
- Understanding Multiple VLAN Registration Protocol (MVRP) on J-EX Series Switches on page 1054
- Understanding Ethernet OAM Connectivity Fault Management for a J-EX Series Switch on page 3463
- Understanding Ethernet OAM Link Fault Management for a J-EX Series Switch on page 3427

- Understanding RSTP for J-EX Series Switches on page 1276
- Understanding STP for J-EX Series Switches on page 1275
- Understanding How to Use sFlow Technology for Network Monitoring on a J-EX Series Switch on page 3283
- Understanding VSTP for J-EX Series Switches on page 1281

[edit routing-instances] Configuration Hierarchy

```
routing-instances routing-instance-name {
    instance-type virtual-router
    interface interface-name
}
```

Related Documentation

- Example: Using Virtual Routing Instances to Route Among VLANs on J-EX Series Switches on page 1112
- Configuring Virtual Routing Instances (CLI Procedure) on page 1142

[edit snmp] Configuration Statement Hierarchy

```
snmp {
    rmon {
        history index {
            bucket-size number;
            interface interface-name;
            interval seconds;
            owner owner-name;
        }
    }
}
```

Related Documentation

- Configuring SNMP (J-Web Procedure) on page 3309
- *Junos OS Network Management Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

[edit virtual-chassis] Configuration Statement Hierarchy

```
virtual-chassis {
    auto-sw-update {
        package-name-edit-virtual-chassis.xml package-name;
    }
    fast-failover (ge | vcp disable | xe);
    id id;
    mac-persistence-timer seconds;
    member member-id {
        mastership-priority number;
        no-management-vlan;
        serial-number;
```

```
    role;
  }
  no-split-detection;
  preprovisioned;
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable> <match
      regex>;
    flag flag ;
  }
}
```

Related Documentation

- Example: Configuring a Virtual Chassis with a Master and Backup in a Single Wiring Closet on page 717
- Example: Configuring a Virtual Chassis Interconnected Across Multiple Wiring Closets on page 733
- Example: Configuring a Virtual Chassis Using a Preprovisioned Configuration File on page 752
- Configuring a Virtual Chassis (CLI Procedure) on page 781
- Configuring a Virtual Chassis (J-Web Procedure) on page 784
- Virtual Chassis Overview on page 691

[edit vlans] Configuration Statement Hierarchy

```
vlans {
  vlan-name {
    description text-description;
    dot1q-tunneling {
      customer-vlans (id | native | range);
      layer2-protocol-tunneling all | protocol-name {
        drop-threshold number;
        shutdown-threshold number;
      }
    }
    filter input filter-name;
    filter output filter-name;
    interface interface-name {
      mapping (native (push | swap) | policy | tag (push | swap));
    }
    l3-interface vlan.logical-interface-number;
    mac-limit number;
    mac-table-aging-time seconds;
    no-local-switching;
    no-mac-learning;
    primary-vlan vlan-name;
    vlan-id number;
    vlan-range vlan-id-low-vlan-id-high;
  }
}
```

Related Documentation

- Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch on page 1063
- Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches on page 1070
- Example: Configure Automatic VLAN Administration Using GVRP on page 1087
- Example: Connecting an Access Switch to a Distribution Switch on page 1078
- Example: Setting Up Q-in-Q Tunneling on J-EX Series Switches on page 1105
- Example: Configuring Layer 2 Protocol Tunneling on J-EX Series Switches on page 1126
- Creating a Private VLAN (CLI Procedure) on page 1143
- Understanding Q-in-Q Tunneling on J-EX Series Switches on page 1051

PART 3

Software Installation

CHAPTER 4

Software Installation Overview

Installation Overview

Understanding Software Installation on J-EX Series Switches

A J-EX Series Switch is delivered with Junos OS preinstalled. As new features and software fixes become available, you must upgrade your software to use them. You can also downgrade Junos OS to a previous release.

This topic covers:

Overview of the Software Installation Process

A J-EX Series switch is delivered with Junos OS preinstalled. When you connect power to the switch, it starts (boots) up from the installed software.

You upgrade Junos OS on a J-EX Series switch by copying a software package to your switch or another system on your local network, then use either the J-Web interface or the CLI to install the new software package on the switch. Finally, you reboot the switch; it boots from the upgraded software. After a successful upgrade, you should back up the new current configuration to a secondary device.

NOTE:
To access the J-Web interface, your management device must have the following software installed:

- Operating system: Microsoft Windows XP Service Pack 3
- Browser version: One of the following. Other browsers might work but are not supported by J-Series platforms.
  - Microsoft Internet Explorer version 7.0
  - Mozilla Firefox version 3.0
- Additional requirements:
  - Only English-language browsers are supported.
  - The browser and the network must be able to receive and process HTTP/1.1 gzip compressed data.

During a successful upgrade, the upgrade package removes all files from **/var/tmp** and completely reinstalls the existing software. It retains configuration files, and similar information, such as secure shell and host keys, from the previous version. The previous software package is preserved in a separate disk partition, and you can manually revert back to it if necessary. If the software installation fails for any reason, such as loss of power during the installation process, the system returns to the originally active installation when you reboot.

Software Package Security

Junos OS is delivered in signed packages that contain digital signatures to ensure it is official software. For more information about signed software packages, see the *Junos OS Installation and Upgrade Guide* at http://www.juniper.net/techpubs/software/junos/.

Installing Software on a Virtual Chassis

You can connect individual J-EX4200 Ethernet Switches together to form one unit and manage the unit as a single chassis, called a Virtual Chassis. The Virtual Chassis operates as a single network entity composed of members. Each member of a Virtual Chassis runs a Junos OS package.

For ease of management, the Virtual Chassis provides flexible methods to upgrade software releases. You can deploy a new software release to all members of a Virtual Chassis or to only a particular member.

Installing Software on J-EX8200 Switches with Redundant Routing Engines

To install software on a J-EX8200 Ethernet Switch that has two Routing Engines with minimal network disruption, you perform a Junos OS installation on each Routing Engine separately, starting with the backup. See “Installing Software on a J-EX8200 Switch with Redundant Routing Engines (CLI Procedure)” on page 71.

Installing Software Using Automatic Software Download

The automatic software download feature uses the DHCP message exchange process to download and install software packages. Users can define a path to a software package on the DHCP server and then the DHCP server communicates this path to J-EX Series switches acting as DHCP clients as part of the DHCP message exchange process. The DHCP clients that have been configured for automatic software download receive these messages and, when the software package name in the DHCP server message is different from that of the software package that booted the DHCP client switch, download and install the software package. See "Upgrading Software Using Automatic Software Download on J-EX Series Switches" on page 82.

Troubleshooting Software Installation

If Junos OS loads but the CLI is not working for any reason, or if the switch has no software installed, you can use the recovery installation procedure to install the software on the switch. See "Troubleshooting Software Installation" on page 91.

NOTE: You can also use this procedure to load two versions of Junos OS in separate partitions on the switch.

Related Documentation

- Downloading Software Packages on page 69
- Installing Software on J-EX Series Switches (J-Web Procedure) on page 75
- Installing Software on a J-EX Series Switch with a Single Routing Engine (CLI Procedure) on page 70
- Installing Software on a J-EX8200 Switch with Redundant Routing Engines (CLI Procedure) on page 71

Junos OS Package Names

You upgrade Junos OS on a J-EX Series Switch by copying a software package to your switch or another system on your local network, then install the new software package on the switch.

A software package name is in the following format:

***package-name-m.nZx.y*-domestic-signed.tgz**

where:

- ***package-name*** is the name of the package—for example, **jinstall-ex-4200**.
- ***m.n*** is the software release, with ***m*** representing the major release number and *n* representing the minor release number—for example, **10.2**.
- ***Z*** indicates the type of software release, where **R** indicates released software and **B** indicates beta-level software.

- ***x.y*** represents the version of the major software release (*x*) and an internal tracking number (*y*)—for example, **1.6**.
- **domestic-signed** is appended to all J-EX Series package names. For most Junos OS packages, **domestic** is used for the United States and Canada and **export** for worldwide distribution. However, for J-EX Series software, **domestic** is used for worldwide distribution as well.

A sample J-EX Series software package name is:

```
jinstall-ex-4200-10.2R1.6-domestic-signed.tgz
```

Related Documentation

- Installing Software on J-EX Series Switches (J-Web Procedure) on page 75
- Installing Software on a J-EX Series Switch with a Single Routing Engine (CLI Procedure) on page 70
- Installing Software on a J-EX8200 Switch with Redundant Routing Engines (CLI Procedure) on page 71
- Downloading Software Packages from Juniper Networks on page 69
- Understanding Software Installation on J-EX Series Switches on page 61

Understanding System Snapshot on J-EX Series Switches

You can create copies of the software running a J-EX Series Switch using the system snapshot feature. The system snapshot feature takes a "snapshot" of the files currently used to run the switch—the complete contents of the **/config** and **/var** directories, which include the running Junos OS, the active configuration, and the rescue configuration—and copies all of these files into an alternate (internal, meaning internal flash, or an external, meaning USB flash) memory source. You can then use this snapshot to boot the switch at the next bootup or as a backup boot option.

You can only use snapshots to move files to external memory if the switch was booted from internal memory, or to move files to internal memory if the switch was booted from external memory. You cannot create a snapshot in the memory source that booted the switch even if the snapshot is being created on a different partition in the same memory source.

Snapshots are particularly useful for moving files onto USB flash drives. You cannot use the **copy** command or any other file-moving technique to move files from an internal memory source to USB memory on the switch.

System snapshots on J-EX Series switches have the following limitations:

- You cannot use snapshots to move files to any destination outside of the switch other than an installed external USB flash drive or to move files between switches that are members of the same virtual chassis.
- Snapshot commands, like other virtual chassis commands, are always executed on a local switch. In cases where a different member switches of the same virtual chassis requests the snapshot, the snapshot command is pushed to the VC member creating the snapshot, executed, and the output is then returned to the switch that initiated the

process. For instance, if the command to create an external snapshot on virtual chassis member 3 is entered from virtual chassis member 1, the snapshot of internal memory on virtual chassis member 3 is taken on external memory on virtual chassis member 3. The output of the process is seen from virtual chassis member 1. No files move between the switches.

Related Documentation

- Understanding Software Installation on J-EX Series Switches on page 61
- Creating a Snapshot and Using It to Boot a J-EX Series Switch on page 80

Licenses Overview

Understanding Software Licenses for the J-EX Series Switch

To enable and use some Junos OS features, you must purchase, install, and manage separate software licenses. The presence on the switch of the appropriate software license keys (“passwords”) determines whether you are eligible to configure and use certain features.

Junos OS feature licenses are device specific. The same feature can be installed and configured on multiple switches. To conform to Junos OS feature licensing requirements, you must purchase a license for each switch.

For a Virtual Chassis deployment, two licenses are recommended for redundancy. These licenses can be based on the serial numbers of any two member switches. If you add additional member switches to the Virtual Chassis configuration, you do not need additional licenses.

Features Requiring a License

The following Junos OS features require an Advanced Feature License (AFL):

- Border Gateway Protocol (BGP) and multiprotocol BGP (MBGP)
- Intermediate System-to-Intermediate System (IS-IS)
- IPv6 routing (except multicast protocols)
- MPLS with RSVP-based label switched paths (LSPs) and MPLS-based circuit cross-connects (CCCs)

You can purchase a license for your J-EX Series switch model. The license allows you to run all the advanced software features on your switch.

For information about how to purchase a software license, contact Dell.

License Warning Messages

For features that require a license, you must install and properly configure a license key to meet the requirements for using the licensable feature. To obtain a license key, use the contact information provided in your Advanced Feature License (AFL) certificate.

If you have not purchased the AFL and installed the license key, you receive warnings after you commit a licensible feature. The system generates system log (**syslog**) alarm messages notifying you that the feature requires a license—for example:

```
Sep  3 05:59:11   craftd[806]:  Minor alarm set, BGP Routing Protocol usage
requires a license
Sep  3 05:59:11   alarmd[805]: Alarm set: License color=YELLOW, class=CHASSIS,
reason=BGP Routing Protocol usage requires a license
Sep  3 05:59:11   alarmd[805]: LICENSE_EXPIRED: License for feature bgp(47) expired
```

Output from the **show system alarms** command displays the active alarms—for example:

```
user@switch> show system alarms
1 alarm currently active
Alarm time               Class  Description
2009-09-03 06:00:11 UTC  Minor  BGP Routing Protocol usage requires a license
```

Every time you edit or view the configuration, a message displays the committed features that require a license. For example, when you edit the BGP configuration, a warning message appears—for example:

```
[edit protocols]
user@switch# bgp
warning: requires 'bgp' license
```

Likewise, viewing the configuration causes the system to display a message—for example:

```
user@switch> show configuration protocols
## Warning: requires 'bgp' license
##
bgp {
    hold-time 10;
    damping;
}
```

Related Documentation

- Managing Licenses for the J-EX Series Switch (CLI Procedure) on page 83
- Managing Licenses for the J-EX Series Switch (J-Web Procedure) on page 84
- Monitoring Licenses for the J-EX Series Switch on page 88
- License Key Components for the J-EX Series Switch on page 66
- J-EX Series Switch Software Features Overview on page 3

License Key Components for the J-EX Series Switch

When you purchase a license for a Junos OS feature that requires a separate license, you receive a license key.

A license key consists of two parts:

- License ID—Alphanumeric string that uniquely identifies the license key. When a license is generated, it is given a license ID.
- License data—Block of binary data that defines and stores all license key objects.

For example, in the following typical license key, the string **JUNOS204558** is the license ID, and the trailing block of data is the license data:

```
JUNOS204558 aeaqea qmijhd amrqha ztfmbu gqzama uqceds
ra32zr lsevik ftvjed o4jy5u fynzzj mgviyl
kgioyf ardb5g sj7wnt rsfked wbjf5a sg
```

The license data defines the device ID for which the license is valid and the version of the license.

Related Documentation

- Managing Licenses for the J-EX Series Switch (CLI Procedure) on page 83
- Managing Licenses for the J-EX Series Switch (J-Web Procedure) on page 84
- Software Licenses for the J-EX Series Switch Overview on page 65

CHAPTER 5

Installing Junos OS

Downloading Software Packages

To upgrade Junos OS on your Dell PowerConnect J-EX Series switch, you can download software packages from the Dell PowerConnect J-Series—Juniper Networks partner website.

Before you can begin to download software upgrades, ensure that you have registered your J-EX Series switch and obtained an account. To register for an account:

1. Locate the chassis serial number (*not* the Dell Service Tag) on your J-EX Series switch.
2. Go to http://www.juniper.net/partners/dell/ and click **Register for an Account**.
3. Fill out the registration information required.

   When your user registration is approved, you receive login information and credentials at the e-mail address you used for registration. If your registration is delayed or additional information is required, you receive a message with further instructions.
4. Save the login information and credentials you receive to use for software download.

To download software upgrades from the Dell PowerConnect J-Series—Juniper Networks partner website:

1. Go to http://www.juniper.net/partners/dell/.
2. Select **J-EX**.
3. Select the appropriate software package for your application. See “Junos OS Package Names” on page 63.
4. Download the software to a local host or to an internal software distribution site.

If you have questions, contact Dell Customer Support at http://www.support.dell.com.

Related Documentation

- Installing Software on J-EX Series Switches (J-Web Procedure) on page 75
- Installing Software on a J-EX Series Switch with a Single Routing Engine (CLI Procedure) on page 70
- Understanding Software Installation on J-EX Series Switches on page 61

Installing Software on a J-EX Series Switch with a Single Routing Engine (CLI Procedure)

You can use this procedure to upgrade Junos OS on a J-EX Series switch with a single Routing Engine, including an individual member of a Virtual Chassis or all members of a Virtual Chassis, or a J-EX8200 switch using a single Routing Engine. To upgrade software on a J-EX8200 switch running two Routing Engines, see "Installing Software on a J-EX8200 Switch with Redundant Routing Engines (CLI Procedure)" on page 71.

To install software upgrades on a J-EX Series switch with a single Routing Engine using the CLI:

1. Download the software package as described in "Downloading Software Packages from Juniper Networks" on page 69.
2. (Optional) Back up the current software configuration to a second storage option. See the *Junos OS Installation and Upgrade Guide* at http://www.juniper.net/techpubs/software/junos/ for instructions on performing this task.
3. (Optional) Copy the software package to the switch. We recommend that you use FTP to copy the file to the **/var/tmp** directory.

   This step is optional because Junos OS can also be upgraded when the software image is stored at a remote location. These instructions describe the software upgrade process for both scenarios.
4. Install the new package on the switch:

   

   NOTE: A reboot, which will occur as part of the execution of the following command, is required to complete the software upgrade. If you want to reboot the switch at a later time, do not use the reboot option at this point of the procedure and enter the request system reboot command at a later time to reboot the switch.

   user@switch> **request system software add *source* reboot**

   Replace ***source*** with one of the following paths:

   - For a software package that is installed from a local directory on the switch—***/pathname/package-name-m.nZx-distribution*.tgz**.
   - For a software package that is downloaded and installed from a remote location:
     - **ftp://*hostname/pathname/package-name-m.nZx-distribution*.tgz**

- **http://*hostname/pathname/package-name-m.nZx-distribution*.tgz**

where *package-name-m.nZx-distribution*.tgz is, for example, **jinstall-ex-4200-10.2R1.8-domestic-signed.tgz**.

Include the optional **member** option to install the software package on only one member of a Virtual Chassis:

user@switch> **request system software add** *source* **member** *member-id* **reboot**

Other members of the Virtual Chassis are not affected. To install the software on all members of the Virtual Chassis, do not include the **member** option.

5. After the reboot has completed, log in and verify that the new version of the software is properly installed:

   user@switch> **show version**

Related Documentation

- Installing Software on J-EX Series Switches (J-Web Procedure) on page 75
- Troubleshooting Software Installation on page 91
- Junos OS Package Names on page 63
- See the *Junos OS System Basics and Services Command Reference* at http://www.juniper.net/techpubs/software/junos/ for details about the **request system software add** command.
- Understanding Software Installation on J-EX Series Switches on page 61

Installing Software on a J-EX8200 Switch with Redundant Routing Engines (CLI Procedure)

For a J-EX8200 switch with redundant Routing Engines, you can minimize disrupting network operation during a Junos OS upgrade by upgrading the Routing Engines separately, starting with the backup Routing Engine.

To upgrade the software package on a J-EX8200 switch with one installed Routing Engine, see "Installing Software on a J-EX Series Switch with a Single Routing Engine (CLI Procedure)" on page 70.

Install the new Junos OS release on the backup Routing Engine while keeping the currently running software version on the master Routing Engine. After making sure that the new software version is running correctly on the backup Routing Engine, switch device control to the backup Routing Engine. Finally, install the new software on the new backup Routing Engine.

To upgrade the Junos OS on the switch, perform the following tasks:

1. Preparing the Switch for the Software Installation on page 72
2. Installing Software on the Backup Routing Engine on page 73
3. Installing Software on the Default Master Routing Engine on page 74
4. Returning Routing Control to the Default Master Routing Engine (Optional) on page 75

Preparing the Switch for the Software Installation

Perform the following steps before installing the software:

1. Log in to the master Routing Engine's console.

   For information on logging in to the Routing Engine through the console port, see "Connecting and Configuring a J-EX Series Switch (CLI Procedure)" on page 161.

2. Enter Junos OS CLI configuration mode:

   a. Start the CLI from the shell prompt:

      ```
      user@switch:RE% cli
      ```

      You will see:

      ```
      {master}
      user@switch>
      ```

   b. Enter configuration mode:

      ```
      user@switch> configure
      ```

      You will see:

      ```
      {master}[[edit]
      user@switch#
      ```

3. Disable GRES (graceful Routing Engine switchover):

   ```
   [edit]
   user@switch# deactivate chassis redundancy graceful-switchover
   ```

4. Save the configuration change on both Routing Engines:

   ```
   [edit]
   user@switch# commit synchronize
   ```

   

   NOTE: To ensure the most recent configuration changes are committed before the software upgrade, perform this step even if GRES was previously disabled.

5. Exit out of the CLI configuration mode:

   ```
   [edit]
   user@switch# exit
   ```

6. (Optional) Back up the current software configuration to a second storage option. See the *Junos OS Installation and Upgrade Guide* at http://www.juniper.net/techpubs/software/junos/ for instructions on performing this task.

Installing Software on the Backup Routing Engine

Once the J-EX8200 switch is ready, you first install the software on the backup Routing Engine. This enables the master Routing Engine to continue operations, minimizing the disruption to your network.

1. Download the software by following the procedures in "Downloading Software Packages from Juniper Networks" on page 69.
2. Copy the software package to the switch. We recommend that you use FTP to copy the file to the **/var/tmp** directory.
3. Log in to the backup Routing Engine's console.
4. Install the new software package:

   user@switch> **request system software add validate**
   /var/tmp/*package-name-m.nZx-distribution*.tgz

   where *package-name-m.nZx-distribution*.tgz is, for example, **jinstall-ex-8200–10.2R1.5–domestic-signed.tgz**.

   For more information on the **request system software add** command, see the *Junos OS System Basics and Services Command Reference* at http://www.juniper.net/techpubs/software/junos/.

   

   NOTE: To abort the installation, do not reboot your device; instead, finish the installation and then issue the **request system software delete** *package-name-m.nZx-distribution*.tgz command, where *package-name-m.nZx-distribution*.tgz is, for example, **jinstall-ex-4200–10.2R1.5–domestic-signed.tgz**. This is your last chance to stop the installation.

5. Reboot to start the new software:

   user@switch> **request system reboot**
   Reboot the system? [yes, no] (no) **yes**

   

   NOTE: You must reboot the switch to load the new installation of Junos OS.

6. After the reboot has completed, log in and verify the new version of the software is properly installed:

   user@switch> **show version**

Installing Software on the Default Master Routing Engine

To switch device control to the backup Routing Engine and then upgrade or downgrade the master Routing Engine software:

1. Log in to the master Routing Engine console port.
2. Transfer device control to the backup Routing Engine:

   user@switch> **request chassis routing-engine master switch**

   

   NOTE: Because GRES is disabled, this switchover causes all line cards in the switch to reload. All network traffic passing through these line cards is lost during the line card reloads.

3. Verify that the default backup Routing Engine (shown as slot 1 in the command output) is now the master Routing Engine:

   user@switch> **show chassis routing-engine**

   You will see:

   ```
   Routing Engine status:
    Slot 0:
      Current state          Backup
      Election priority      Master (default)
   Routing Engine status:
    Slot 1:
      Current state          Master
      Election priority      Backup (default)
   ```

4. Install the new software package using the **request system software add** command:

   user@switch> **request system software add validate**
   /var/tmp/jinstall-ex-8200-10.2R1.5-domestic-signed.tgz

5. Reboot the Routing Engine:

   user@switch> **request system reboot**
   Reboot the system? [yes, no] (no) **yes**

   When the reboot completes, the prompt will reappear. Wait for this prompt to reappear before proceeding to the next step.

6. Log in to the default backup Routing Engine (slot 1) through the console port.
7. Re-enable GRES:

   [edit]
   user@switch# **activate chassis redundancy graceful-switchover**

   Re-enabling GRES allows any future Routing Engine switchovers to occur without the loss of any network traffic.

8. Enter the **commit synchronize** command to save the configuration change:

   [edit]

```
user@switch# commit synchronize
```

9. Log in and verify the version of the software installed.

   If you want to return routing control to the Routing Engine that was the master Routing Engine at the beginning of the procedure (the default master Routing Engine), perform the next task.

Returning Routing Control to the Default Master Routing Engine (Optional)

The switch can maintain normal operations with the Routing Engine in slot 1 acting as the master Routing Engine after the software upgrade, so only perform this task if you want to return routing control to the default master Routing Engine in slot 0.

1. Transfer routing control back to the default master Routing Engine:

   ```
   user@switch> request chassis routing-engine master switch
   ```

2. Verify that the default master Routing Engine (slot 0) is indeed the master Routing Engine:

   ```
   user@switch> show chassis routing-engine
   ```

   You will see:

   ```
   Routing Engine status:
    Slot 0:
      Current state          Master
      Election priority      Master (default)
   Routing Engine status:
    Slot 1:
      Current state          Backup
      Election priority      Backup (default)
   ```

Related Documentation

- Installing Software on J-EX Series Switches (J-Web Procedure) on page 75
- Troubleshooting Software Installation on page 91
- Junos OS Package Names on page 63
- Understanding Software Installation on J-EX Series Switches on page 61
- Understanding J-EX8208 Switch Component and Functionality Redundancy

Installing Software on J-EX Series Switches (J-Web Procedure)

You can upgrade software packages on a single fixed-configuration switch, on an individual member of a Virtual Chassis, or for all members of a Virtual Chassis.

You can use the J-Web interface to install software upgrades from a server using FTP or HTTP, or by copying the file to the J-EX Series switch.

NOTE:
To access the J-Web interface, your management device must have the following software installed:

- Operating system: Microsoft Windows XP Service Pack 3
- Browser version: One of the following. Other browsers might work but are not supported by J-Series platforms.
  - Microsoft Internet Explorer version 7.0
  - Mozilla Firefox version 3.0
- Additional requirements:
  - Only English-language browsers are supported.
  - The browser and the network must be able to receive and process HTTP/1.1 gzip compressed data.

This topic describes:

Installing Software Upgrades from a Server

To install software upgrades from a remote server by using FTP or HTTP:

1. Download the software package as described in "Downloading Software Packages from Juniper Networks" on page 69.
2. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
3. In the J-Web interface, select **Maintain>Software>Install Package**.
4. On the Install Remote page, enter information into the fields described in Table 22 on page 76.
5. Click **Fetch and Install Package**. The software is activated after the switch has rebooted.

Table 22: Install Remote Summary

Field	Function	Your Action
Package Location (required)	Specifies the FTP or HTTP server, file path, and software package name.	Type the full address of the software package location on the FTP or HTTP server—one of the following: ftp://*hostname/pathname/package-name* http://*hostname/pathname/package-name*

Table 22: Install Remote Summary *(continued)*

Field	Function	Your Action
User	Specifies the username, if the server requires one.	Type the username.
Password	Specifies the password, if the server requires one.	Type the password.
Reboot If Required	If this box is checked, the switching platform is automatically rebooted when the upgrade is complete.	Check the box if you want the switching platform to reboot automatically when the upgrade is complete.

Installing Software Upgrades by Uploading Files

To install software upgrades by uploading files:

1. Download the software package.
2. In the J-Web interface, select **Maintain>Software>Upload Package**.
3. On the Upload Package page, enter information into the fields described in Table 23 on page 77.
4. Click **Upload and Install Package**. The software is activated after the switching platform has rebooted.

Table 23: Upload Package Summary

Field	Function	Your Action
File to Upload (required)	Specifies the location of the software package.	Type the location of the software package, or click **Browse** to navigate to the location.
Reboot If Required	Specifies that the switching platform is automatically rebooted when the upgrade is complete.	Select the check box if you want the switching platform to reboot automatically when the upgrade is complete.

Related Documentation

- Installing Software on a J-EX Series Switch with a Single Routing Engine (CLI Procedure) on page 70
- Understanding Software Installation on J-EX Series Switches on page 61
- Troubleshooting Software Installation on page 91

Rebooting or Halting the J-EX Series Switch (J-Web Procedure)

You can use the J-Web interface to schedule a reboot or to halt the switching platform.

To reboot or halt the switching platform by using the J-Web interface:

1. In the J-Web interface, select **Maintain>Reboot**.
2. Select one:
   - **Reboot Immediately**—Reboots the switching platform immediately.
   - **Reboot in *number of* minutes**—Reboots the switch in the number of minutes from now that you specify.
   - **Reboot when the system time is *hour:minute*** —Reboots the switch at the absolute time that you specify, on the current day. You must select a 2-digit hour in 24-hour format and a 2-digit minute.
   - **Halt Immediately**— Stops the switching platform software immediately. After the switching platform software has stopped, you can access the switching platform through the console port only.
3. (Optional) In the Message box, type a message to be displayed to any users on the switching platform before the reboot occurs.
4. Click **Schedule**. The J-Web interface requests confirmation to perform the reboot or halt.
5. Click **OK** to confirm the operation.
   - If the reboot is scheduled to occur immediately, the switch reboots. You cannot access the J-Web interface until the switch has restarted and the boot sequence is complete. After the reboot is complete, refresh the browser window to display the J-Web interface login page.
   - If the reboot is scheduled to occur in the future, the Reboot page displays the time until reboot. You have the option to cancel the request by clicking **Cancel Reboot** on the J-Web interface Reboot page.
   - If the switch is halted, all software processes stop and you can access the switching platform through the console port only. Reboot the switch by pressing any key on the keyboard.

Related Documentation

- Starting the J-Web Interface on page 136

CHAPTER 6

Booting the Switch, Upgrading Software, and Managing Licenses

Booting the Switch

Booting a J-EX Series Switch Using a Software Package Stored on a USB Flash Drive

There are two methods of getting Junos OS onto a USB flash drive before using the software to boot the switch. You can pre-install the software onto the USB flash drive before inserting the USB flash drive into the USB port, or you can use the system snapshot feature to copy files from internal switch memory to the USB flash drive.

To move files into USB flash memory using a system snapshot and use those files to boot the switch, see “Creating a Snapshot and Using It to Boot a J-EX Series Switch” on page 80. We recommend that you use this method to boot the switch from a USB flash drive if your switch is running properly.

If you need to pre-install the software onto the USB flash drive, you can use the method described in this topic. Pre-installing Junos OS onto a USB flash drive to boot the switch can be done at any time and is particularly useful when the switch boots to the loader prompt because the switch cannot locate Junos OS in internal flash memory.

Ensure that you have the following tools and parts available to boot the switch from a USB flash drive:

- A USB flash drive that meets the J-EX Series switch USB port specifications. See USB Port Specifications for a J-EX Series Switch.
- A computer or other device that you can use to download the software package from the Internet and copy it to the USB flash drive.

To download a Junos OS package onto a USB flash drive before inserting the USB flash drive:

1. Download the Junos OS package that you would like to place onto the J-EX Series switch from the Internet onto the USB flash drive using your computer or other device. See "Downloading Software Packages from Juniper Networks" on page 69.
2. Remove the USB flash drive from the computer or other device.
3. Insert the USB flash drive into the USB port on the switch.
4. This step can only be performed when the prompt for the loader script (**loader>**) is displayed. The loader script starts when Junos OS loads but the CLI is not working for any reason or if the switch has no software installed.

   Install the software package onto the switch:

   loader> **install** ***source***

   where ***source*** represents the name and location of the Junos OS package on the USB flash drive. The Junos OS package on a flash drive is commonly stored in the root drive as the only file—for example, **file:///jinstall-ex-4200-10.2R1.5-domestic-signed.tgz**.

Related Documentation

- Installing Software on a J-EX Series Switch with a Single Routing Engine (CLI Procedure) on page 70
- Installing Software on J-EX Series Switches (J-Web Procedure) on page 75
- See Rear Panel of a J-EX4200 Switch for USB port location.
- See Switch Fabric and Routing Engine (SRE) Module in a J-EX8208 Switch for USB port location.
- See Routing Engine (RE) Module in a J-EX8216 Switch for USB port location.
- Understanding Software Installation on J-EX Series Switches on page 61

Creating a Snapshot and Using It to Boot a J-EX Series Switch

The system snapshot feature takes a "snapshot" of the files currently used to run the J-EX Series switch—the complete contents of the **/config** and **/var** directories, which include the running Junos OS, the active configuration, and the rescue configuration—and copies all of these files into an alternate (internal, meaning internal flash, or an external, meaning USB flash) memory source. You can then use these snapshots to boot the switch at the next bootup or as a backup boot option.

This topic includes the following tasks:

1. Creating a Snapshot on a USB Flash Drive and Using It to Boot the Switch on page 80
2. Creating a Snapshot on an Internal Flash Drive and Using it to Boot the Switch on page 81

Creating a Snapshot on a USB Flash Drive and Using It to Boot the Switch

A snapshot can be created on USB flash memory after a switch is booted using files stored in internal memory.

Ensure that you have the following tools and parts available before creating a snapshot on a USB Flash drive:

- A USB flash drive that meets the J-EX Series switch USB port specifications. See USB Port Specifications for a J-EX Series Switch.

To create a snapshot on USB flash memory and use it to boot the switch:

1. Place the snapshot into USB flash memory:

   user@switch> **request system snapshot partition media external slice 1**

   NOTE: This example uses the partition option. If you have already created a partition for the snapshot, you don't need to use the partition option.

2. (Optional) Perform this step if you want to boot the switch now using the snapshot stored on the USB flash drive. If you created the snapshot as a backup, do not perform this step.

   - To reboot the switch using the most recently created snapshot:

     user@switch> **request system reboot media external**

   - To reboot the switch using a snapshot in a specific partition on the USB flash drive:

     user@switch> **request system reboot media external slice 1**

Creating a Snapshot on an Internal Flash Drive and Using it to Boot the Switch

A snapshot can be created on internal memory after a switch is booted using files stored in external memory.

To create a snapshot in internal memory and use it to boot the switch:

1. Place the snapshot files in internal memory:

   user@switch> **request system snapshot parition media internal slice 1**

   NOTE: This example uses the partition option. If you have already created a partition for the snapshot, you don't need to use the partition option.

2. (Optional) Perform this step if you want to boot the switch now using the newly created snapshot. If you created the snapshot as a backup, do not perform this step.

   - To reboot the switch using the most recently created snapshot:

     user@switch> **request system reboot media internal**

   - To reboot the switch using a snapshot in a specific partition in internal memory:

     user@switch> **request system reboot media internal slice 1**

Related Documentation

- Verifying That a System Snapshot Was Created on a J-EX Series Switch on page 88

- Understanding System Snapshot on J-EX Series Switches on page 64

Upgrading Software

- Upgrading Software Using Automatic Software Download on J-EX Series Switches on page 82

Upgrading Software Using Automatic Software Download on J-EX Series Switches

The automatic software download feature uses the DHCP message exchange process to download and install software packages. You configure the automatic software download feature on J-EX Series switches acting as DHCP clients. You must enable automatic software download on the J-EX Series switch before the software upgrade can occur.

You configure a path to a software package file on the DHCP server. The server communicates the path to the software package file through DHCP server messages.

If you enable automatic software download, the DHCP client J-EX Series switch compares the software package name in the DHCP server message to the name of the software package that booted the switch. If the software packages are different, the DHCP client J-EX Series switch downloads and installs the software package specified in the DHCP server message.

Before you upgrade software using automatic software download, ensure that you have configured DHCP services for the switch, including configuring a path to a boot server and a boot file. See the *Junos OS System Basics Configuration Guide* at http://www.juniper.net/techpubs/software/junos/ for information about using the CLI to configure DHCP services and settings. See “Configuring DHCP Services (J-Web Procedure)” on page 447 for information about using the J-Web interface to configure DHCP services and settings.

To enable automatic software download on a J-EX Series switch acting as a DHCP client:

```
[edit chassis]
user@switch# set auto-image-upgrade
```

Once automatic software download is enabled on your DHCP client J-EX Series switch and once DHCP services are enabled on your network, an automatic software download can occur at any time as part of the DHCP message exchange process.

If an automatic software download occurs, you see the following message on the switch:

```
Auto-image upgrade started
On successful installation system will reboot automatically
```

The switch reboots automatically to complete the upgrade.

Related Documentation

- Verifying That Automatic Software Download Is Working Correctly on page 87
- Understanding Software Installation on J-EX Series Switches on page 61
- DHCP Services for J-EX Series Switches Overview on page 445

Managing Licenses

Managing Licenses for the J-EX Series Switch (CLI Procedure)

To enable and use some Junos OS features on a J-EX Series switch, you must purchase, install, and manage separate software licenses. Each switch requires one license. For a Virtual Chassis deployment, two licenses are recommended for redundancy. After you have configured the features, you see a warning message if the switch does not have a license for the feature.

Before you begin managing licenses, be sure that you have:

- Obtained the needed licenses. For information about how to purchase software licenses, contact Dell.
- Understand what makes up a license key. For more information, see “License Key Components for the J-EX Series Switch” on page 66.

This topic includes the following tasks:

Adding New Licenses

To add one or more new license keys on the switch, with the CLI:

1. Add the license key or keys:
   - To add one or more license keys from a file or URL, specify the filename of the file or the URL where the key is located:

     ```
     user@switch> request system license add filename | url
     ```

   - To add a license key from the terminal:

     ```
     user@switch> request system license add terminal
     ```

2. When prompted, enter the license key, separating multiple license keys with a blank line.

   If the license key you enter is invalid, an error appears in the CLI output when you press Ctrl+d to exit the license entry mode.

Deleting Licenses

To delete one or more license keys from the switch with the CLI, specify the license ID:

```
user@switch> request system license delete license-id
```

You can delete only one license at a time.

Saving License Keys

To save the installed license keys to a file (which can be a URL) or to the terminal:

```
user@switch> request system license save filename | url
```

For example, the following command saves the installed license keys to a file named **license.conf**:

```
user@switch> request system license save ftp://user@switch/license.conf
```

Related Documentation

- Managing Licenses for the J-EX Series Switch (J-Web Procedure) on page 84
- Monitoring Licenses for the J-EX Series Switch on page 88
- Understanding Software Licenses for the J-EX Series Switch on page 65

Managing Licenses for the J-EX Series Switch (J-Web Procedure)

To enable and use some Junos OS features on a J-EX Series switch, you must purchase, install, and manage separate software licenses. Each switch requires one license. For a Virtual Chassis deployment, two licenses are recommended for redundancy. After you have configured the features, you see a warning message if the switch does not have a license for the feature.

Before you begin managing licenses, be sure that you have:

- Obtained the needed licenses. For information about how to purchase software licenses, contact Dell.
- Understand what makes up a license key. For more information, see “License Key Components for the J-EX Series Switch” on page 66.

NOTE:
To access the J-Web interface, your management device must have the following software installed:

- Operating system: Microsoft Windows XP Service Pack 3
- Browser version: One of the following. Other browsers might work but are not supported by J-Series platforms.
  - Microsoft Internet Explorer version 7.0
  - Mozilla Firefox version 3.0
- Additional requirements:
  - Only English-language browsers are supported.
  - The browser and the network must be able to receive and process HTTP/1.1 gzip compressed data.

This topic includes the following tasks:

Adding New Licenses

To add one or more new license keys on the switch, with the J-Web license manager:

1. In the J-Web interface, select **Maintain>Licenses**.
2. Under Installed Licenses, click **Add** to add a new license key or keys.
3. Do *one* of the following, using a blank line to separate multiple license keys:
   - In the License File URL box, type the full URL to the destination file containing the license key or keys to be added.
   - In the License Key Text box, paste the license key text, in plain-text format, for the license to be added.
4. Click **OK** to add the license key or keys.

A list of features that use the license key is displayed. The table also lists the ID, state, and version of the license key.

Deleting Licenses

To delete one or more license keys from a switch with the J-Web license manager:

1. In the J-Web interface, select **Maintain>Licenses**.
2. Select the check box of the license or licenses you want to delete.

3. Click **Delete**.

Displaying License Keys

To display the license keys installed on a switch with the J-Web license manager:

1. In the J-Web interface, select **Maintain>Licenses**.
2. Under Installed Licenses, click **Display Keys** to display all the license keys installed on the switch.

A screen displaying the license keys in text format appears. Multiple licenses are separated by a blank line.

Downloading Licenses

To download the license keys installed on the switch with the J-Web license manager:

1. In the J-Web interface, select **Maintain>Licenses**.
2. Under Installed Licenses, click **Download Keys** to download all the license keys installed on the switch to a single file.
3. Select **Save it to disk** and specify the file to which the license keys are to be written. You can also download the license file to your system.

Related Documentation

- Managing Licenses for the J-EX Series Switch (CLI Procedure) on page 83
- Monitoring Licenses for the J-EX Series Switch on page 88
- Understanding Software Licenses for the J-EX Series Switch on page 65

CHAPTER 7

Verifying Software Installation

Routine Monitoring

Verifying That Automatic Software Download Is Working Correctly

Purpose Verify that the automatic software download feature is working correctly.

Action Use the **show system services dhcp client *interface-name*** command to verify that the automatic software download feature has been used to install a software package.

```
user@switch> show system services dhcp client ge-0/0/1.0
 Logical Interface Name       ge-0/0/1.0
  Hardware address        00:0a:12:00:12:12
  Client Status           bound
  Vendor Identifier       ether
  Server Address          10.1.1.1
  Address obtained        10.1.1.89
  Lease Obtained at       2009-08-20 18:13:04 PST
  Lease Expires at        2009-08-22 18:13:04 PST

DHCP Options :
  Name: name-server, Value: [ 10.209.194.131, 2.2.2.2, 3.3.3.3 ]
  Name: server-identifier, Value: 10.1.1.1
  Name: router, Value: [ 10.1.1.80 ]
  Name: boot-image,
  Value: jinstall-ex-4200-10.2R1.5-domestic-signed.tgz
  Name: boot-image-location,
  Value: 10.1.1.25:/bootfiles/
```

Meaning The output from this command shows the name and location of the software package under **DHCP options** when automatic software download was last used to install a software package. The sample output in **DHCP options** shows that the last DHCP server message to arrive on the DHCP client had a boot server address of **192.168.1.165** and a boot file named **jinstall-ex-4200-10.2R1.5-domestic-signed.tgz**. If automatic software

download was enabled on this client switch during the last DHCP message exchange, these values were used by the switch to upgrade the software.

Related Documentation

- Upgrading Software Using Automatic Software Download on J-EX Series Switches on page 82
- DHCP Services for J-EX Series Switches Overview on page 445

Verifying That a System Snapshot Was Created on a J-EX Series Switch

Purpose Verify that a system snapshot was created with the proper files on a J-EX Series switch.

Action View the snapshot:

```
user@switch> show system snapshot media external
Information for snapshot on external (da1s1)
Creation date: Oct 1320:23:23 2009
Junos version on snapshot:
jbase : 10.0I20090726_0011_user
jcrypto-ex: 10.0I20090726_0011_user
jdocs-ex: 10.0I20090726_0011_user
jkernel-ex: 10.0I20090726_0011_user
jroute-ex: 10.0I20090726_0011_user
jswitch-ex: 10.0I20090726_0011_user
jweb-ex: 10.0I20090726_0011_user
jpfe-ex42x: 10.0I20090726_0011_user
```

Meaning The output shows the date and time when the snapshot was created and the packages that are part of the snapshot. The date and time match the time when you created the snapshot.

You can compare the output of this command to the output of the **show system software** command to ensure that the snapshot contains the same packages as the software currently running the switch.

Related Documentation

- Creating a Snapshot and Using It to Boot a J-EX Series Switch on page 80

Monitoring Licenses

- Monitoring Licenses for the J-EX Series Switch on page 88

Monitoring Licenses for the J-EX Series Switch

To enable and use some Junos OS features on the J-EX Series switch, you must purchase, install, and manage the appropriate software licenses. Each switch requires one license. For a Virtual Chassis deployment, two licenses are recommended for redundancy.

To monitor your installed licenses, perform the following tasks:

- Displaying Installed Licenses and License Usage Details on page 89
- Displaying Installed License Keys on page 90

Displaying Installed Licenses and License Usage Details

Purpose Verify that the expected license is installed and active on the switch and fully covers the switch configuration.

Action From the CLI, enter the **show system license** command. (To display only the **License usage** list, enter the **show system license usage** command. To display only the **Licenses installed** output, enter **show system license installed**.)

```
user@switch> show system license
License usage:

                              Licenses     Licenses    Licenses    Expiry
  Feature name                     used    installed      needed
  bgp                                 1            1           0    permanent
  isis                                0            1           0    permanent
  ospf3                               0            1           0    permanent
  ripng                               0            1           0    permanent
  mpls                                0            1           0    permanent


Licenses installed:

  License identifier: JUNOS204558

  License version: 2

  Valid for device: BN0208380000

  Features:

    ex-series - Licensed routing protocols in ex-series

      permanent
```

Meaning The output shows the license or licenses (for Virtual Chassis deployments) installed on the switch and license usage. Verify the following information:

- If a feature that requires a license is configured (used), a license is installed on the switch. The **Licenses needed** column must show that no licenses are required.
- The appropriate number of licenses is installed. Each switch requires one license. For a Virtual Chassis deployment, two licenses are recommended for redundancy.
- The expected license is installed.

Displaying Installed License Keys

Purpose Verify that the expected license keys are installed on the switch.

Action From the CLI, enter the **show system license keys** command.

```
user@switch> show system license keys
JUNOS204558 aeaqea qmijhd amrqha ztfmbu gqzama uqceds

            ra32zr lsevik ftvjed o4jy5u fynzzj mgviyl

            kgioyf ardb5g sj7wnf rsdked wbjf5a sg
```

Meaning The output shows the license key or keys (for Virtual Chassis deployments) installed on the switch. Verify that each expected license key is present.

Related Documentation

- Managing Licenses for the J-EX Series Switch (CLI Procedure) on page 83
- Managing Licenses for the J-EX Series Switch (J-Web Procedure) on page 84
- Understanding Software Licenses for the J-EX Series Switch on page 65

CHAPTER 8

Troubleshooting Software Installation

Troubleshooting Software Installation

Recovering from a Failed Software Upgrade on a J-EX Series Switch

Problem If Junos OS loads but the CLI is not working for any reason, or if the switch has no software installed, you can use this recovery installation procedure to install Junos OS.

Solution If there is already a Junos OS image on the system, you can install the new Junos OS package in a separate partition and both images will remain on the system, or you can wipe the disk clean before the new installation proceeds.

If there is no Junos OS image on the system, follow the instructions in "Booting a J-EX Series Switch Using a Software Package Stored on a USB Flash Drive" on page 79 to get an image on the system and boot the switch.

To perform a recovery installation:

1. Power on the switch. The loader script starts.

   After the message **Loading /boot/defaults/loader.conf** displays, you are prompted with:

   Hit [Enter] to boot immediately, or space bar for command prompt.

2. Press the space bar to enter the manual loader. The **loader>** prompt displays.

3. Enter the following command:

   loader> **install [– –format] [– –external]** ***source***

   where:

   - **format**—Use this option to wipe the installation media before installing the software package. If you do not include this option, the system installs the new Junos OS package in a different partition from that of the most recently installed Junos OS package.

- **external**—Use this option to install the software package onto an external media.
- **source**—Represents the name and location of the Junos OS package either on a server on the network or as a file on the USB flash drive:
  - Network address of the server and the path on the server; for example, **tftp://192.17.1.28/junos/jinstall-ex-4200-10.2R1.5-domestic-signed.tgz**
  - The Junos OS package on a USB device is commonly stored in the root drive as the only file; for example, **file:///jinstall-ex-4200-10.2R1.5-domestic-signed.tgz**

The boot process proceeds as normal and ends with a login prompt.

Rebooting from the Inactive Partition

Problem A J-EX Series switch ships with Junos OS loaded on the system disk in partition 1. The first time you upgrade, the new software package is installed in partition 2. When you finish the installation and reboot, partition 2 becomes the active partition. Similarly, subsequent software packages are installed in the inactive partition which becomes the active partition when you reboot at the end of the installation process.

If you performed an upgrade and rebooted, the system resets the active partition. You can use this procedure to manually boot from the inactive partition.

NOTE: If you have completed the installation of the software image but have not yet rebooted, you can issue the request system software rollback command to return to the original software installation package.

Solution Reboot from the inactive partition:

```
user@switch> request system reboot partition alternate
```

NOTE: If you cannot access the CLI, you can reboot from the inactive partition using the following procedure from the loader script prompt:

1. Unload and clear the interrupted boot from the active partition:

   ```
   loader> unload
   loader> unset vfs.root.mountfrom
   ```

2. Select the new (inactive) partition to boot from:

   ```
   loader> set currdev=disk media s partition:
   ```

   where *media* is either 0 (internal) or 1 (external) and *partition* indicates the partition number, either 1 or 2.

   You must include the colon (:) at the end of this command.

3. Boot Junos OS from the inactive partition:

   ```
   loader> boot
   ```

Related Documentation

- Installing Software on a J-EX Series Switch with a Single Routing Engine (CLI Procedure) on page 70
- Installing Software on J-EX Series Switches (J-Web Procedure) on page 75
- Understanding Software Installation on J-EX Series Switches on page 61

CHAPTER 9

Configuration Statements for Software Installation

[edit chassis] Configuration Statement Hierarchy

```
chassis {
  aggregated-devices {
    ethernet {
      device-count number;
    }
  }
  auto-image-upgrade;
  fpc slot {
    pic pic-number {
      sfpplus {
        pic-modemode;
      }
    }
    power-budget-priority priority;
  }
  lcd-menu fpc slot-number {
    menu-item (menu-name | menu-option);
  }
  psu {
    redundancy {
      n-plus-n;
    }
  }
  redundancy {
    graceful-switchover;
  }
}
```

Related Documentation

- Upgrading Software Using Automatic Software Download on J-EX Series Switches on page 82
- Configuring the LCD Panel on J-EX Series Switches (CLI Procedure) on page 166
- Configuring Graceful Routing Engine Switchover in a Virtual Chassis Configuration (CLI Procedure) on page 801

- Configuring Power Supply Redundancy (CLI Procedure) on page 307
- Configuring the Power Priority of Line Cards (CLI Procedure) on page 308
- Configuring Nonstop Software Upgrade (CLI Procedure)

auto-image-upgrade

Syntax auto-image-upgrade;

Hierarchy Level [edit chassis]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable automatic software download on a J-EX Series switch acting as a DHCP client.

The DHCP client J-EX Series switch compares the software package name in the DHCP server message to the name of the software package that booted the switch. If the software packages are different, the DHCP client J-EX Series switch downloads and installs the software package specified in the DHCP server message.

Before you upgrade software using automatic software download, ensure that you have configured DHCP services for the switch, including configuring a path to a boot server and a boot file. See the *Junos OS System Basics Configuration Guide* at http://www.juniper.net/techpubs/software/junos/index.html for information about using the CLI to configure DHCP services and settings. See "Configuring DHCP Services (J-Web Procedure)" on page 447 for information about using the J-Web interface to configure DHCP services and settings.

Default Automatic software download is disabled.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Upgrading Software Using Automatic Software Download on J-EX Series Switches on page 82
- Understanding Software Installation on J-EX Series Switches on page 61
- DHCP Services for J-EX Series Switches Overview on page 445

CHAPTER 10

Operational Mode Commands for Software Installation

request system license add

Syntax request system license add (*filename* | terminal)

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Add a license key.

Options *filename*—License key from a file or URL. Specify the filename or the URL where the key is located.

terminal—License key from the terminal.

Required Privilege Level maintenance

List of Sample Output **request system license add on page 98**

Output Fields When you enter this command, you are provided feedback on the status of your request.

request system license add

```
user@host> request system license add terminal
```

request system license delete

Syntax request system license delete *license-id*

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Delete a license key. You can delete only one license at a time.

Options *license-id*—License ID that uniquely identifies a license key.

Required Privilege Level maintenance

List of Sample Output **request system license delete on page 99**

Output Fields When you enter this command, you are provided feedback on the status of your request.

request system license delete

```
user@host> request system license delete G03000002223
```

request system license save

Syntax request system license save (*filename* | terminal)

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Save installed license keys to a file or URL.

Options *filename*—License key from a file or URL. Specify the filename or the URL where the key is located.

terminal—License key from the terminal.

Required Privilege Level maintenance

List of Sample Output **request system license save on page 100**

Output Fields When you enter this command, you are provided feedback on the status of your request.

request system license save

```
user@host> request system license save ftp://user@host/license.conf
```

request system reboot

Syntax

request system reboot
<other-routing-engine>
<at *time*>
<in *minutes*>
<media (compact-flash | disk | removable-compact-flash | usb)>
<message "*text*">

Syntax (J-EX Series Switch)

request system reboot
<all-members>
<at *time*>
<in *minutes*>
<local>
<media (external | internal)>
<member *member-id*>
<message "*text*">
<other-routing-engine>
<slice *slice*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Reboot the software.

Options

none—Reboot the software immediately.

all-members—(J-EX4200 switches only) (Optional) Reboot all members of the Virtual Chassis configuration.

at *time*—(Optional) Time at which to reboot the software, specified in one of the following ways:

- **now**—Stop or reboot the software immediately. This is the default.
- ***+minutes***—Number of minutes from now to reboot the software.
- ***yymmddhhmm***—Absolute time at which to reboot the software, specified as year, month, day, hour, and minute.
- ***hh:mm***—Absolute time on the current day at which to stop the software, specified in 24-hour time.

in *minutes*—(Optional) Number of minutes from now to reboot the software. This option is an alias for the **at *+minutes*** option.

local—(J-EX4200 switches only) (Optional) Reboot the local Virtual Chassis member.

media (compact-flash | disk | removable-compact-flash | usb)—(Optional) Boot medium for next boot.

media (external | internal)—(J-EX Series switches only) (Optional) Reboot the boot media:

- **external**—Reboot the external mass storage device.
- **internal**—Reboot the internal flash device.

member *member-id*—(J-EX4200 switches only) (Optional) Reboot the specified member of the Virtual Chassis configuration Replace ***member-id*** with a value from 0 through 9.

message "*text*"—(Optional) Message to display to all system users before stopping or rebooting the software.

other-routing-engine—(Optional) Reboot the other Routing Engine from which the command is issued. For example, if you issue the command from the master Routing Engine, the backup Routing Engine is rebooted. Similarly, if you issue the command from the backup Routing Engine, the master Routing Engine is rebooted.

slice *slice*—(J-EX Series switches only) (Optional) Reboot a partition on the boot media. This option has the following suboptions:

- **1**—Power off partition 1.
- **2**—Power off partition 2.
- **alternate**—Reboot from the alternate partition.

Additional Information

Reboot requests are recorded in the system log files, which you can view with the **show log** command (see **show log**). Also, the names of any running processes that are scheduled to be shut down are changed. You can view the process names with the **show system processes** command (see **show system processes**).

NOTE: To reboot a router that has two Routing Engines, reboot the backup Routing Engine (if you have upgraded it) first, and then reboot the master Routing Engine.

Required Privilege Level

maintenance

Related Documentation

- **clear system reboot on page 204**

List of Sample Output

request system reboot on page 102
request system reboot (at 2300) on page 102
request system reboot (in 2 Hours) on page 103
request system reboot (Immediately) on page 103
request system reboot (at 1:20 AM) on page 103

Output Fields

When you enter this command, you are provided feedback on the status of your request.

request system reboot

```
user@host> request system reboot
Reboot the system ? [yes,no] (no)
```

request system reboot (at 2300)

```
user@host> request system reboot at 2300 message ?Maintenance time!?
Reboot the system ? [yes,no] (no) yes

shutdown: [pid 186]
```

```
*** System shutdown message from root@berry.network.net ***
System going down at 23:00
```

request system reboot (in 2 Hours)

The following example, which assumes that the time is 5 PM (17:00), illustrates three different ways to request the system to reboot in two hours:

```
user@host> request system reboot at +120
user@host> request system reboot in 120
user@host> request system reboot at 19:00
```

request system reboot (Immediately)

```
user@host> request system reboot at now
```

request system reboot (at 1:20 AM)

To reboot the system at 1:20 AM, enter the following command. Because 1:20 AM is the next day, you must specify the absolute time.

```
user@host> request system reboot at 06060120
request system reboot at 120
Reboot the system at 120? [yes,no] (no) yes
```

request system reboot

Syntax

request system reboot
<all-members | local | member *member-id*>
<at *time*>
<in *minutes*>
<media (external | internal)>
<message "*text*">
<slice (1 | 2 | alternate)>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Reboot Junos OS.

Reboot requests are recorded in the system log files, which you can view with the **show log** command. You can view the process names with the **show system processes** command.

Options

none—Reboots the software immediately.

all-members | **local** | **member** ***member-id***—(J-EX4200 switch only) (Optional) Specify which member of the Virtual Chassis to reboot:

- **all-members**—Reboots each switch that is a member of the Virtual Chassis.
- **local**—Reboots the local switch, meaning the switch you are logged into, only.
- **member** ***member-id***—Reboots the specified member switch of the Virtual Chassis.

at *time*—(Optional) Time at which to reboot the software, specified in one of the following ways:

- ***+minutes***—Number of minutes from now to reboot the software.
- ***hh:mm***—Absolute time on the current day at which to reboot the software, specified in 24-hour time.
- **now**—Stop or reboot the software immediately. This is the default.
- ***yymmddhhmm***—Absolute time at which to reboot the software, specified as year, month, day, hour, and minute.

in *minutes*—(Optional) Number of minutes from now to reboot the software. This option is an alias for the **at** ***+minutes*** option.

media (external | internal)—(Optional) Boot medium for the next boot. The external option reboots the switch using a software package stored on an external boot source, such as a USB flash drive. The internal option reboots the switch using a software package stored in an internal memory source.

message "*text*"—(Optional) Message to display to all system users before rebooting the software.

slice (1 | 2 | alternate)—(Optional) Reboot using the specified partition on the boot media. This option has the following suboptions:

- **1**—Reboot from partition 1.
- **2**—Reboot from partition 2.
- **alternate**—Reboot from the alternate partition, which is the partition that did not boot the switch at the last bootup.

Required Privilege Level maintenance

Related Documentation
- **clear system reboot on page 204**

Output Fields When you enter this command, you are provided feedback on the status of your request.

request system reboot

```
user@host> request system reboot
Reboot the system ? [yes,no] (no)
```

request system reboot (at 2300)

```
user@host> request system reboot at 2300 message ?Maintenance time!?
Reboot the system ? [yes,no] (no) yes

shutdown: [pid 186]
*** System shutdown message from root@berry.network.net ***
System going down at 23:00
```

request system reboot (in 2 Hours) The following example, which assumes that the time is 5 PM (17:00), illustrates three different ways to request the system to reboot in two hours:

```
user@host> request system reboot at +120
user@host> request system reboot in 120
user@host> request system reboot at 19:00
```

request system reboot (Immediately)

```
user@host> request system reboot at now
```

request system reboot (at 1:20 AM) To reboot the system at 1:20 AM, enter the following command. Because 1:20 AM is the next day, you must specify the absolute time.

```
user@host> request system reboot at 06060120
request system reboot at 120
Reboot the system at 120? [yes,no] (no) yes
```

request system snapshot

Syntax

request system snapshot
<as-primary>
<all-members | local | member *member-id*>
<media (external | internal)>
<partition>
<re0 | re1 | routing-engine *routing-engine-id*>
<slice (1 | 2 | alternate)>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Take a snapshot of the files currently used to run the switch—the complete contents of the **/config** and **/var** directories, which include the running Junos OS, the active configuration, and the rescue configuration—and copy all of these files into an alternate (internal, meaning internal flash, or an external, meaning USB flash) memory source.

Options

none—Create a snapshot on the alternate media, meaning the external media if you booted the switch using software stored on internal media or internal media if you booted the switch using software stored on external media.

all-members | **local** | **member** *member-id*—(J-EX4200 switch only) (Optional) Specify where to place the snapshot in Virtual Chassis configurations:

- **all-members**—Create a snapshot for each switch that is a member of the Virtual Chassis.
- **local**—Create a snapshot on the local switch only.
- **member** *member-id*—Create a snapshot for the specified member or member switches of the Virtual Chassis.

as-primary—(Optional) Create a bootable snapshot.

NOTE: The snapshot is always bootable on J-EX Series switches. The as-primary option has no effect on snapshots on J-EX Series switches.

media (**external** | **internal**)—(Optional) Specify the destination media location for the snapshot. The **external** option copies the snapshot to an external mass storage device, such as a USB flash drive. The **internal** option copies the snapshot to an internal memory source, such as internal flash memory.

partition—(Optional) Partition the destination media before copying over the snapshot.

re0 | **re1** | **routing-engine** ***routing-engine-id***—(J-EX8200 switch only) Specify where to place the snapshot in dual Routing Engine configurations.

- **re0**—Create a snapshot on Routing Engine 0.
- **re1**—Create a snapshot on Routing Engine 1.

- **routing-engine***routing-engine-id*—Create a snapshot on the specified Routing Engine.

slice (**1** | **2** | **alternate**)—(Optional) Specify the destination partition for the snapshot:

- **1**—Copy the snapshot to partition 1.
- **2**—Copy the snapshot to partition 2.
- **alternate**—Copy the snapshot to the alternate partition, which is the partition that did not boot the switch at the last bootup.

Required Privilege Level view

Related Documentation

- **show system snapshot on page 122**
- Creating a Snapshot and Using It to Boot a J-EX Series Switch on page 80

Output Fields When you enter this command, you are provided feedback on the status of your request.

request system snapshot media external slice 1

```
user@switch> request system snapshot media external slice 1
```

request system software add

Syntax request system software add *package-name*
<best-effort-load>
<delay-restart>
<force>
<no-copy>
<no-validate>
<re0 | re1>
<reboot>
<unlink>
<validate>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Install a software package or bundle on the router or switch.

Options *package-name*—Location from which the software package or bundle is to be installed. For example:

- **/var/tmp/*package-name***—For a software package or bundle that is being installed from a local directory on the router or switch.
- ***protocol*://*hostname*/*pathname*/*package-name***—For a software package or bundle that is to be downloaded and installed from a remote location. Replace ***protocol*** with one of the following:
  - **ftp**—File Transfer Protocol.
    Use **ftp://*hostname*/*pathname*/*package-name***. To specify authentication credentials, use **ftp://<*username*>:<*password*>@*hostname*/*pathname*/*package-name***. To have the system prompt you for the password, specify **prompt** in place of the password. If a password is required, and you do not specify the password or **prompt**, an error message is displayed.
  - **http**—Hypertext Transfer Protocol.
    Use **http://*hostname*/*pathname*/*package-name***. To specify authentication credentials, use **http://<*username*>:<*password*>@*hostname*/*pathname*/*package-name***. If a password is required and you omit it, you are prompted for it.
  - **scp**—Secure copy (available only for Canada and U.S. version).
    Use **scp://*hostname*/*pathname*/*package-name***. To specify authentication credentials, use **scp://<*username*>:<*password*>@*hostname*/*pathname*/*package-name***.

NOTE: The *pathname* in the protocol is the relative path to the user's home directory on the remote system and not the root directory.

best-effort-load—(Optional) Activate a partial load and treat parsing errors as warnings instead of errors.

delay-restart—(Optional) Install software package or bundle, but do not restart software processes.

force—(Optional) Force the addition of the software package or bundle (ignore warnings).

no-copy—(Optional) Install a software package or bundle, but do not save copies of package or bundle files.

no-validate—(Optional) When loading a software package or bundle with a different release, suppress the default behavior of the **validate** option.

re0 | re1—(Optional) On routers that support dual or redundant Routing Engines, load a software package or bundle on the Routing Engine in slot 0 (**re0**) or Routing Engine in slot 1 (**re1**).

reboot—(Optional) After adding the software package or bundle, reboot the system.

unlink—(Optional) Remove the software package from this directory after a successful upgrade is completed.

validate—(Optional) Validate the software package or bundle against the current configuration as a prerequisite to adding the software package or bundle. This is the default behavior when the software package or bundle being added is a different release.

Additional Information

Before upgrading the software on the router or switch, when you have a known stable system, issue the **request system snapshot** command to back up the software, including the configuration, to the **/altroot** and **/altconfig** file systems. After you have upgraded the software on the router or switch and are satisfied that the new package or bundle is successfully installed and running, issue the **request system snapshot** command again to back up the new software to the **/altroot** and **/altconfig** file systems.

After you run the **request system snapshot** command, you cannot return to the previous version of the software, because the running and backup copies of the software are identical.

If you are upgrading more than one package at the same time, delete the operating system package, **jkernel**, last. Add the operating system package, **jkernel**, first and the routing software package, **jroute**, last. If you are upgrading all packages at once, delete and add them in the following order:

```
user@host> request system software add /var/tmp/jbase
user@host> request system software add /var/tmp/jkernel
user@host> request system software add /var/tmp/jpfe
user@host> request system software add /var/tmp/jdocs
user@host> request system software add /var/tmp/jroute
user@host> request system software add /var/tmp/jcrypto
```

Required Privilege Level maintenance

Related Documentation

- request system software delete on page 111
- request system software rollback on page 113
- request system storage cleanup on page 227

List of Sample Output request system software add validate on page 110

Output Fields When you enter this command, you are provided feedback on the status of your request.

request system software add validate

```
user@host> request system software add validate /var/tmp/ jinstall-7.2R1.7-domestic-signed.tgz
Checking compatibility with configuration
Initializing...
Using jbase-7.1R2.2
Using /var/tmp/jinstall-7.2R1.7-domestic-signed.tgz
Verified jinstall-7.2R1.7-domestic.tgz signed by PackageProduction_7_2_0
Using /var/validate/tmp/jinstall-signed/jinstall-7.2R1.7-domestic.tgz
Using /var/validate/tmp/jinstall/jbundle-7.2R1.7-domestic.tgz
Checking jbundle requirements on /
Using /var/validate/tmp/jbundle/jbase-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jkernel-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jcrypto-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jpfe-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jdocs-7.2R1.7.tgz
Using /var/validate/tmp/jbundle/jroute-7.2R1.7.tgz
Validating against /config/juniper.conf.gz
mgd: commit complete
Validation succeeded
Validating against /config/rescue.conf.gz
mgd: commit complete
Validation succeeded
Installing package '/var/tmp/jinstall-7.2R1.7-domestic-signed.tgz' ...
Verified jinstall-7.2R1.7-domestic.tgz signed by PackageProduction_7_2_0
Adding jinstall...

WARNING:     This package will load JUNOS 7.2R1.7 software.
WARNING:     It will save JUNOS configuration files, and SSH keys
WARNING:     (if configured), but erase all other files and information
WARNING:     stored on this machine.  It will attempt to preserve dumps
WARNING:     and log files, but this can not be guaranteed.  This is the
WARNING:     pre-installation stage and all the software is loaded when
WARNING:     you reboot the system.

Saving the config files ...
Installing the bootstrap installer ...

WARNING:     A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:     'request system reboot' command when software installation is
WARNING:     complete. To abort the installation, do not reboot your system,
WARNING:     instead use the 'request system software delete jinstall'
WARNING:     command as soon as this operation completes.

Saving package file in /var/sw/pkg/jinstall-7.2R1.7-domestic-signed.tgz ...
Saving state for rollback ...
```

request system software delete

Syntax request system software delete *software-package*
<force>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Remove a software package or bundle from the router or switch.

CAUTION: Before removing a software package or bundle, make sure that you have already placed the new software package or bundle that you intend to load onto the router or switch.

Options *software-package*—Software package or bundle name. You can delete any or all of the following software bundles or packages:

- **jbase**—(Optional) Junos OS base software suite
- **jcrypto**—(Optional, in domestic version only) Junos OS security software
- **jdocs**—(Optional) Junos OS online documentation file
- **jkernel**—(Optional) Junos OS kernel software suite
- **jpfe**—(Optional) Junos OS Packet Forwarding Engine support
- **jroute**—(Optional) Junos OS routing software suite
- **junos**—(Optional) Junos OS base software

force—(Optional) Ignore warnings and force removal of the software.

Additional Information Before upgrading the software on the router or switch, when you have a known stable system, issue the **request system snapshot** command to back up the software, including the configuration, to the **/altroot** and **/altconfig** file systems. After you have upgraded the software on the router or switch and are satisfied that the new packages are successfully installed and running, issue the **request system snapshot** command again to back up the new software to the **/altroot** and **/altconfig** file systems. After you run the **request system snapshot** command, you cannot return to the previous version of the software, because the running and backup copies of the software are identical.

Required Privilege Level maintenance

Related Documentation

- **request system software add on page 108**
- **request system software rollback on page 113**
- **request system software validate on page 115**

List of Sample Output **request system software delete jdocs on page 112**

Output Fields When you enter this command, you are provided feedback on the status of your request.

request system software delete jdocs The following example displays the system software packages before and after the **jdocs** package is deleted through the **request system software delete** command:

```
user@host> show system software
Information for jbase:

Comment:
JUNOS Base OS Software Suite [7.2R1.7]


Information for jcrypto:

Comment:
JUNOS Crypto Software Suite [7.2R1.7]


Information for jdocs:

Comment:
JUNOS Online Documentation [7.2R1.7]


Information for jkernel:

Comment:
JUNOS Kernel Software Suite [7.2R1.7]

...

user@host> request system software delete jdocs
Removing package 'jdocs' ...


user@host> show system software
Information for jbase:

Comment:
JUNOS Base OS Software Suite [7.2R1.7]


Information for jcrypto:

Comment:
JUNOS Crypto Software Suite [7.2R1.7]


Information for jkernel:

Comment:
JUNOS Kernel Software Suite [7.2R1.7]

...
```

request system software rollback

Syntax

request system software rollback

Syntax (J-EX Series Switch)

request system software rollback
<all-members>
<local>
<member *member-id*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Revert to the software that was loaded at the last successful **request system software add** command.

Options

none—Revert to the set of software as of the last successful **request system software add**.

all-members—(J-EX4200 switches only) (Optional) Attempt to roll back to the previous set of packages on all members of the Virtual Chassis configuration.

local—(J-EX4200 switches only) (Optional) Attempt to roll back to the previous set of packages on the local Virtual Chassis member.

member *member-id*—(J-EX4200 switches only) (Optional) Attempt to roll back to the previous set of packages on the specified member of the Virtual Chassis configuration. Replace ***member-id*** with a value from 0 through 9.

Additional Information

Use this command only to recover from a failed software upgrade—you cannot issue this command to return to the previously installed software after using a **jinstall** package. To return to the previously installed software, use the corresponding **jinstall** package.

A software rollback fails if any required package (or a **jbundle** package containing the required package) cannot be found in **/var/sw/pkg**.

Required Privilege Level

maintenance

Related Documentation

- **request system software add on page 108**
- **request system software delete on page 111**
- **request system software validate on page 115**
- **request system configuration rescue delete on page 376**
- **request system configuration rescue save on page 377**

List of Sample Output

request system software rollback on page 114

Output Fields

When you enter this command, you are provided feedback on the status of your request.

request system software rollback

```
user@host> request system software rollback
Verified SHA1 checksum of ./jbase-7.2R1.7.tgz
Verified SHA1 checksum of ./jdocs-7.2R1.7.tgz
Verified SHA1 checksum of ./jroute-7.2R1.7.tgz
Installing package './jbase-7.2R1.7.tgz' ...
Available space: 35495 require: 7335
Installing package './jdocs-7.2R1.7.tgz' ...
Available space: 35339 require: 3497
Installing package './jroute-7.2R1.7.tgz' ...
Available space: 35238 require: 6976
NOTICE: uncommitted changes have been saved in
/var/db/config/juniper.conf.pre-install
Reloading /config/juniper.conf.gz ...
Activating /config/juniper.conf.gz ...
mgd: commit complete
Restarting mgd ...
Restarting aprobed ...
Restarting apsd ...
Restarting cosd ...
Restarting fsad ...
Restarting fud ...
Restarting gcdrd ...
Restarting ilmid ...
Restarting irsd ...
Restarting l2tpd ...
Restarting mib2d ...
Restarting nasd ...
Restarting pppoed ...
Restarting rdd ...
Restarting rmopd ...
Restarting rtspd ...
Restarting sampled ...
Restarting serviced ...
Restarting snmpd ...
Restarting spd ...
Restarting vrrpd ...

WARNING: cli has been replaced by an updated version:
CLI release 7.2R1.7 built by builder on 2005-04-22 02:03:44 UTC
Restart cli using the new version ? [yes,no] (yes) yes

Restarting cli ...
user@host
```

request system software validate

Syntax request system software validate *package-name*

Syntax (J-EX Series Switch) request system software validate
<member *member-id*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Validate candidate software against the current configuration of the router or switch.

Options member *member-id*—(J-EX4200 switches only) (Optional) Validate the software bundle or package on the specified member of the Virtual Chassis configuration. Replace ***member-id*** with a value from 0 through 9.

package-name—Name of the software bundle or package to test.

Required Privilege Level maintenance

Related Documentation

- **request system software add on page 108**
- **request system software delete on page 111**
- **request system software rollback on page 113**

List of Sample Output

request system software validate (Successful Case) on page 116
request system software validate (Failure Case) on page 116

Output Fields When you enter this command, you are provided feedback on the status of your request.

request system software validate (Successful Case)

```
user@host> request system software validate /var/sw/pkg/jbundle-5.3I20020124_0520_sjg.tgz
Checking compatibility with configuration
Initializing...
Using /packages/jbase-5.3I20020122_1901_sjg
Using /var/sw/pkg/jbundle-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jbase-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jkernel-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jcrypto-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jpfe-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jdocs-5.3I20020124_0520_sjg.tgz
Using /var/chroot/var/tmp/jbundle/jroute-5.3I20020124_0520_sjg.tgz
Validating against /config/juniper.conf.gz
mgd: commit complete

WARNING: cli has been replaced by an updated version:
CLI release 5.3I0 built by sjg on 2002-01-24 05:23:53 UTC
Restart cli using the new version ? [yes,no] (yes)
```

request system software validate (Failure Case)

```
user@host> request system software validate 6.3/
Pushing bundle to lcc0-re0
error: Failed to transfer package to lcc0-re0


user@host> request system software validate test
Pushing bundle to lcc0-re0
Pushing bundle to lcc2-re0

lcc0-re0:
gzip: stdin: not in gzip format
tar: child returned status 1
ERROR: Not a valid package: /var/tmp/test
```

show system autoinstallation status

Syntax show system autoinstallation status

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display autoinstallation status information.

Options This command has no options.

Required Privilege Level view

List of Sample Output **show system autoinstallation status on page 117**

show system autoinstallation status

```
user@host> show system autoinstallation status
Autoinstallation status:
Master state: Active
Last committed file: None
Configuration server of last committed file: 0.0.0.0
Interface:
  Name: fe-0/0/1
  State: None
  Address acquisition:
    Protocol: DHCP Client
    Acquired address: None
    Protocol: RARP Client
    Acquired address: None
```

show system boot-messages

Syntax show system boot-messages

Syntax (J-EX Series Switch)
show system boot-messages
<all-members>
<local>
<member *member-id*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display initial messages generated by the system kernel upon startup. These messages are the contents of **/var/run/dmesg.boot**.

Options none—Display all boot time messages.

all-members—(J-EX4200 switches only) (Optional) Display boot time messages on all members of the Virtual Chassis configuration.

local—(J-EX4200 switches only) (Optional) Display boot time messages on the local Virtual Chassis member.

member *member-id*—(J-EX4200 switches only) (Optional) Display boot time messages on the specified member of the Virtual Chassis configuration. Replace ***member-id*** with a value from 0 through 9.

Required Privilege Level view

show system license

Syntax show system license
<installed | keys | usage>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display licenses and information about how they are used.

Options none—Display all license information.

installed—(Optional) Display installed licenses only.

keys—(Optional) Display a list of license keys. Use this information to verify that each expected license key is present.

usage—(Optional) Display the state of licensed features.

Required Privilege Level maintenance

List of Sample Output
show system license on page 120
show system license installed on page 120
show system license keys on page 120
show system license usage on page 120

Output Fields Table 24 on page 119 lists the output fields for the **show system license** command. Output fields are listed in the approximate order in which they appear.

Table 24: show system license Output Fields

Field Name	Field Description
Feature name	Name assigned to the configured feature. You use this information to verify that all the features for which you installed licenses are present.
Licenses used	Number of licenses used by a router or switch. You use this information to verify that the number of licenses used matches the number configured. If a licensed feature is configured, the feature is considered used.
Licenses installed	Information about the installed license key: • **License identifier**—Identifier associated with a license key. • **State**—State of the license key:**valid** or **invalid**. An **invalid** state indicates that the key was entered incorrectly or is not valid for the specific device. • **License version**—Version of a license. The version indicates how the license is validated, the type of signature, and the signer of the license key. • **Valid for device**—Device that can use a license key. • **Group defined**—Group membership of a device. • **Features**—Feature associated with a license, such as data link switching (DLSw).
Licenses needed	Number of licenses required for features being used but not yet properly licensed.

Table 24: show system license Output Fields *(continued)*

Field Name	Field Description
Expiry	Amount of time left within the grace period before a license is required for a feature being used.

show system license

```
user@host> show system license

License usage:
                                Licenses     Licenses    Licenses    Expiry
  Feature name                       used    installed      needed
  subscriber-accounting                 0            1           0    permanent
  subscriber-authentication             0            1           0    permanent
  subscriber-address-assignment         0            1           0    permanent
  subscriber-vlan                       0            1           0    permanent
  subscriber-ip                         0            1           0    permanent
  scale-subscriber                      0         1000           0    permanent
  scale-l2tp                            0         1000           0    permanent
  scale-mobile-ip                       0         1000           0    permanent

Licenses installed:
  License identifier: XXXXXXXXXX
  License version: 2
  Features:
    subscriber-accounting - Per Subscriber Radius Accounting
      permanent
    subscriber-authentication - Per Subscriber Radius Authentication
      permanent
    subscriber-address-assignment - Radius/SRC Address Pool Assignment
      permanent
    subscriber-vlan  - Dynamic Auto-sensed Vlan
      permanent
    subscriber-ip    - Dynamic and Static IP
      permanent
```

show system license installed

```
user@host> show system license installed
  License identifier: XXXXXXXXXX
  License version: 2
  Features:
    subscriber-accounting - Per Subscriber Radius Accounting
      permanent
    subscriber-authentication - Per Subscriber Radius Authentication
      permanent
    subscriber-address-assignment - Radius/SRC Address Pool Assignment
      permanent
    subscriber-vlan  - Dynamic Auto-sensed Vlan
      permanent
    subscriber-ip    - Dynamic and Static IP
      permanent
```

show system license keys

```
user@host> show system license keys
XXXXXXXXXX xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx
           xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx xxxxxx
           xxxxxx xxxxxx xxx
```

show system license usage

```
user@host> show system license usage
                                Licenses     Licenses    Licenses    Expiry
  Feature name                       used    installed      needed
```

```
subscriber-accounting              1          1          0    permanent
subscriber-authentication          1          1          0    permanent
subscriber-address-assignment      1          1          0    permanent
subscriber-vlan                    0          1          0    permanent
subscriber-ip                      0          1          0    permanent
```

show system snapshot

Syntax

show system snapshot
<all-members | local | member *member-id*>
<media (external | internal)>
<slice (1 | 2 | alternate)>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display the complete collection of files in a snapshot.

Options

none—Display the system snapshot on the alternate media, which is the media that does not have the software packages that last booted the switch.

all-members | local | member *member-id*—(J-EX4200 switch only) Display the snapshot in a Virtual Chassis configuration:

- **all-members**—Display the snapshot for each switch that is a member of the Virtual Chassis.
- **local**—Display the snapshot on the switch that you are currently logged into.
- **member *member-id***—Display the snapshot for the specified member switch of the Virtual Chassis.

media (external | internal)—(Optional) Display the destination media location for the snapshot. The **external** option specifies the snapshot on an external mass storage device, such as a USB flash drive. The **internal** option specifies the snapshot on an internal memory source, such as internal flash memory.

slice (1 | 2 | alternate)—Display the snapshot in a partition:

- **1**—Display the snapshot in partition 1.
- **2**—Display the snapshot in partition 2.
- **alternate**—Display the snapshot in the alternate partition, which is the partition that did not boot the switch at the last bootup.

Required Privilege Level

view

Related Documentation

- **request system snapshot on page 106**
- Creating a Snapshot and Using It to Boot a J-EX Series Switch on page 80
- Verifying That a System Snapshot Was Created on a J-EX Series Switch on page 88

show system snapshot media external

```
user@switch> show system snapshot media external
Information for snapshot on external (da1s1)
Creation date: Oct 13 20:23:23 2009
JUNOS version on snapshot:
jbase  : 10.0I20090726_0011_user
jcrypto-ex: 10.0I20090726_0011_user
```

```
jdocs-ex: 10.0I20090726_0011_user
jkernel-ex: 10.0I20090726_0011_user
jroute-ex: 10.0I20090726_0011_user
jswitch-ex: 10.0I20090726_0011_user
jweb-ex: 10.0I20090726_0011_user
jpfe-ex42x: 10.0I20090726_0011_user
```

PART 4

User Interfaces

CHAPTER 11

User Interfaces Overview

User Interfaces—Overview

CLI User Interface Overview

You can use two interfaces to monitor, configure, troubleshoot, and manage a J-EX Series Switch: the J-Web graphical user interface and the Junos OS command-line interface (CLI). Both of these user interfaces are shipped with the switch. This topic describes the CLI. For information about the J-Web user interface, see "J-Web User Interface for J-EX Series Switches Overview" on page 129.

CLI Overview

Junos OS CLI is a command shell that runs on top of a UNIX-based operating system kernel. The CLI provides command help and command completion.

The CLI also provides a variety of UNIX utilities, such as Emacs-style keyboard sequences that allow you to move around on a command line and scroll through recently executed commands, regular expression matching to locate and replace values and identifiers in a configuration, filter command output, or log file entries, store and archive router files on a UNIX-based file system, and exit from the CLI environment and create a UNIX C shell or Bourne shell to navigate the file system, manage switch processes, and so on.

CLI Help and Command Completion

To access CLI Help, type a question mark (**?**) at any level of the hierarchy. The system displays a list of the available commands or statements and a short description of each.

To complete a command, statement, or option that you have partially typed, press the Tab key or the Spacebar. If the partially typed letters uniquely identify a command, the complete command name appears. Otherwise, a beep indicates that you have entered an ambiguous command and the possible completions are displayed. This completion feature also applies to other strings, such as filenames, interface names, usernames, and configuration statements.

CLI Command Modes

The CLI has two modes, operational mode and configuration mode.

In operational mode, you enter commands to monitor and troubleshoot switch hardware and software and network connectivity. Operational mode is indicated by the > prompt—for example, **user@switch>**.

In configuration mode, you can define all properties of Junos OS, including interfaces, VLANs, Virtual Chassis information, routing protocols, user access, and several system hardware properties.

To enter configuration mode, enter the **configure** command: .

```
user@switch> configure
```

Configuration mode is indicated by the **#** prompt, and includes the current location in the configuration hierarchy—for example:

```
[edit interfaces ge-0/0/12]
user@switch#
```

In configuration mode, you are actually viewing and changing the candidate configuration file. The candidate configuration allows you to make configuration changes without causing operational changes to the current operating configuration, called the active configuration. When you commit the changes you added to the candidate configuration, the system updates the active configuration. Candidate configurations enable you to alter your configuration without causing potential damage to your current network operations.

To activate your configuration changes, enter the **commit** command.

To return to operational mode, go to the top of the configuration hierarchy and then quit—for example:

```
[edit interfaces ge-0/0/12]
user@switch# top
[edit]
user@switch# exit
```

You can also activate your configuration changes and exit configuration mode with a single command, **commit and-quit**. This command succeeds only if there are no mistakes or syntax errors in the configuration.

TIP: When you commit the candidate configuration, you can require an explicit confirmation for the commit to become permanent by using the commit

confirmed command. This is useful for verifying that a configuration change works correctly and does not prevent management access to the switch. After you issue the commit confirmed command, you must issue another commit command within the defined period of time (10 minutes by default) or the system reverts to the previous configuration.

Related Documentation

- J-EX Series Switch Software Features Overview on page 3
- *Junos OS CLI User Guide* at http://www.juniper.net/techpubs/software/junos/.

J-Web User Interface for J-EX Series Switches Overview

You can use two interfaces to monitor, configure, troubleshoot, and manage a J-EX Series Switch: the J-Web graphical user interface and the Junos OS command-line interface (CLI). Both of these user interfaces are shipped with the switch. This topic describes the J-Web interface. You can navigate the J-Web interface, scroll pages, and expand and collapse elements as you do in a typical Web browser interface. For information about the CLI user interface, see "CLI User Interface Overview" on page 127.

NOTE:

To access the J-Web interface, your management device must have the following software installed:

- Operating system: Microsoft Windows XP Service Pack 3
- Browser version: One of the following. Other browsers might work but are not supported by J-Series platforms.
  - Microsoft Internet Explorer version 7.0
  - Mozilla Firefox version 3.0
- Additional requirements:
  - Only English-language browsers are supported.
  - The browser and the network must be able to receive and process HTTP/1.1 gzip compressed data.

Each page of the J-Web interface is divided into panes.

- Top pane—Displays system identity information and links.
- Main pane—Location where you monitor, configure, diagnose (troubleshoot), and manage (maintain) the switch by entering information in text boxes, making selections, and clicking buttons.
- Side pane—Displays suboptions of the Monitor, Configure, Troubleshoot, or Maintain task currently displayed in the main pane. Click a suboption to access it in the main pane.

The layout of the panes allows you to quickly navigate through the interface. Table 25 on page 130 summarizes the elements of the J-Web interface.

The J-Web interface provides CLI tools that allow you to perform all of the tasks that you can perform from the Junos OS command-line interface (CLI), including a CLI Viewer to view the current configuration, a CLI Editor for viewing and modifying the configuration, and a Point & Click CLI editor that allows you to click through all of the available CLI statements.

Table 25: J-Web Interface

J-Web Interface Element	Description
Top Pane	
Host	The hostname of the switch.
Logged in as: username	The user name you used to log in to the switch.

Commit Options	A set of options using which you can configure committing multiple changes with a single commit.
	• Commit—Commits the candidate configuration of the current user session, along with changes from other user sessions.
	• Compare—Displays the XML log of pending configurations on the device.
	• Discard—Discards the candidate configuration of the current user session, along with changes from other user sessions.
	• Preference—Indicates your choice of committing all configurations changes together or committing each configuration change immediately. The two commit options are:
	  ◦ **Commit changes immediately**—Sets the system to force an immediate commit on every page after every configuration change.
	  ◦ **Validate changes until explicit commit**—Loads all configuration changes for an accumulated single commit. If there are errors in loading the configuration, the errors are logged. This is the default mode.
	NOTE: There are some pages on which configuration changes must be committed immediately. For such pages, if you configure the commit options for a single commit, the system displays warning notifications that remind you to commit your changes immediately. An example for such a page is **Switching**.

Help	Displays links to information on help and the J-Web interface. • Help Contents—View context-sensitive help topics. • About—Displays information about the J-Web interface, such as the version number.
Logout	Ends your current login session with the switch and returns you to the login page.

Table 25: J-Web Interface *(continued)*

J-Web Interface Element	Description
Taskbar	Menu of J-Web main options. Click the tab to access an option. • Dashboard—Displays a high-level, graphical view of the chassis and status of the switch. It displays system health information, alarms, and system status. • Configure—Configure the switch, and view configuration history. • Monitor—View information about configuration and hardware on the switch. • Maintain—Manage files and licenses, upgrade software, and reboot the switch. • Troubleshoot—Run diagnostic tools to troubleshoot network issues.
Main Pane	
Help (?) icon	Displays useful information—such as the definition, format, and valid range of an option—when you move the cursor over the question mark.
Red asterisk (*)	Indicates a required field.
Icon legend	(Applies to the Point & Click CLI editor only) Explains icons that appear in the user interface to provide information about configuration statements: • C—Comment. Move your cursor over the icon to view a comment about the configuration statement. • I—Inactive. The configuration statement does not affect the switch. • M—Modified. The configuration statement has been added or modified. • *—Mandatory. The configuration statement must have a value.
Task Pane	
Configuration hierarchy	(Applies to the Junos OS CLI configuration editor only) Displays the hierarchy of committed statements in the switch configuration. • Click **Expand all** to display the entire hierarchy. • Click **Hide all** to display only the statements at the top level. • Click plus signs (+) to expand individual items. • Click minus signs (-) to hide individual items.

Related Documentation

- Using the Commit Options to Commit Configuration Changes (J-Web Procedure) on page 334
- J-EX Series Switch Software Features Overview on page 3
- J-EX4200 Switches Hardware Overview on page 25
- J-EX Series Switch Software Features Overview on page 3
- Connecting and Configuring a J-EX Series Switch (J-Web Procedure) on page 163
- CLI User Interface Overview on page 127

Understanding J-Web Configuration Tools

The J-Web graphical user interface (GUI) allows you to monitor, configure, troubleshoot, and manage the switching platform by means of a Web browser with Hypertext Transfer

Protocol (HTTP) or HTTP over Secure Sockets Layer (HTTPS) enabled. The J-Web interface provides access to all the configuration statements supported by the switch, so you can fully configure the switch without using the CLI.

The J-Web interface provides three methods of configuring the switch:

- Configure menu
- Point & Click CLI Editor
- CLI Editor

Table 26 on page 132 gives a comparison of the three methods of configuration.

Table 26: Switching Platform Configuration Interfaces

Tool	Description	Function	Use
Configure menu	Web browser pages for setting up the switch quickly and easily without configuring each statement individually. For example, use the Virtual Chassis Configuration page to configure the Virtual Chassis parameters on the switch.	Configure basic switch platform services: • Interfaces • Switching • Virtual Chassis • Security • Services • System Properties • Routing	Use for basic configuration.
Point & Click CLI editor	Web browser pages divided into panes in which you can do any of the following: • Expand the entire configuration hierarchy and click a configuration statement to view or edit. The main pane displays all the options for the statement, with a text box for each option. • Paste a complete configuration hierarchy into a scrollable text box, or edit individual lines. • Upload or download a complete configuration. • Roll back to a previous configuration. • Create or delete a rescue configuration.	Configure all switching platform services: • System parameters • User Accounting and Access • Interfaces • VLAN properties • Virtual Chassis properties • Secure Access • Services • Routing protocols	Use for complete configuration if you are not familiar with the Junos OS CLI or prefer a graphical interface.
CLI editor	Interface in which you do any of the following: • Type commands on a line and press **Enter** to create a hierarchy of configuration statements. • Create an ASCII text file that contains the statement hierarchy. • Upload a complete configuration, or roll back to a previous configuration. • Create or delete a rescue configuration.	Configure all switching platform services: • System parameters • User Accounting and Access • Interfaces • VLAN properties • Virtual Chassis properties • Secure Access • Services • Routing protocols	Use for complete configuration if you know the Junos OS CLI or prefer a command interface.

Related Documentation

- Understanding J-Web User Interface Sessions on page 133
- J-Web User Interface for J-EX Series Switches Overview on page 129
- Connecting and Configuring a J-EX Series Switch (J-Web Procedure) on page 163
- Configuration Files Terms on page 322

Understanding J-Web User Interface Sessions

You establish a J-Web session with the switch through an HTTP-enabled or HTTPS-enabled Web browser. The HTTPS protocol, which uses 128-bit encryption, is available only in domestic versions of Junos OS. To use HTTPS, you must have installed a certificate on the switch and enabled HTTPS. See "Generating SSL Certificates to Be Used for Secure Web Access" on page 398.

When you attempt to log in through the J-Web interface, the switch authenticates your username with the same methods used for Telnet and SSH.

If the switch does not detect any activity through the J-Web interface for 15 minutes, the session times out and is terminated. You must log in again to begin a new session.

To explicitly terminate a J-Web session at any time, click **Logout** in the top pane.

Related Documentation

- J-Web User Interface for J-EX Series Switches Overview on page 129
- Configuring Management Access for the J-EX Series Switch (J-Web Procedure) on page 395

CHAPTER 12

Using the Configuration Tools

Using the CLI Terminal

The J-Web CLI terminal provides access to the Junos OS command line interface (CLI) through the J-Web interface. The functionality and behavior of the CLI available through the CLI terminal page is the same as that of the Junos OS CLI available through the switch console. The CLI terminal supports all CLI commands and other features such as CLI help and autocompletion. Using the CLI terminal page you can fully configure, monitor, and manage the switch.

To access the J-Web interface, your management device must have the following software installed:

- Operating system: Microsoft Windows XP Service Pack 3
- Browser version: One of the following. Other browsers might work but are not supported by J-Series platforms.
  - Microsoft Internet Explorer version 7.0
  - Mozilla Firefox version 3.0
- Additional requirements:
  - Only English-language browsers are supported.
  - The browser and the network must be able to receive and process HTTP/1.1 gzip compressed data.
- Before you can use the CLI terminal, you must configure the domain name and hostname of the switch. See “Configuring System Identity for the J-EX Series Switch (J-Web Procedure)” on page 168 for more information.
- To access the CLI through the J-Web interface, your management device requires the following features:
  - SSH access—Enable Secure shell (SSH) on your system. SSH provides a secured method of logging in to the switch, to encrypt traffic so that it is not intercepted. If SSH is not enabled on the system, the CLI terminal page displays an error.

- Java applet support—Make sure that your Web browser supports Java applets.
- JRE installed on the client—Install Java Runtime Environment (JRE) version 1.4 or later on your system. JRE is a software package that must be installed on a system to run Java applications. Download the latest JRE version from the Java Software website **http://www.java.com/**. Installing JRE installs Java plug-ins, which once installed, load automatically and transparently to render Java applets.

NOTE: The CLI terminal is supported on JRE version 1.4 and later only.

To access the CLI terminal, select **Troubleshoot >CLI Terminal**.

Related Documentation

- CLI User Interface Overview on page 127
- Understanding J-Web Configuration Tools on page 131

Starting the J-Web Interface

You can use the J-Web graphical interface to configure and manage the J-EX Series switch.

To start the J-Web interface:

1. Launch your HTTP-enabled or HTTPS-enabled Web browser.

   To use HTTPS, you must have installed a certificate on the switch and enabled HTTPS.

2. After **http://** or **https://** in your Web browser, type the hostname or IP address of the switch and press **Enter**.

   The J-Web login page appears.

3. On the login page, type your username and password, and click **Log In**.

   To correct or change the username or password you typed, click **Reset**, type the new entry or entries, and click **Log In**.

   

   NOTE: The default username is root with no password. You must change this during initial configuration or the system does not accept the configuration.

   The Chassis Dashboard information page appears.

To explicitly terminate a J-Web session at any time, click **Logout** in the top pane.

Related Documentation

- J-Web User Interface for J-EX Series Switches Overview on page 129
- Understanding How to Use the J-Web Interface to View System Information

CHAPTER 13

Operational Mode Commands for User Interfaces

set cli complete-on-space

Syntax set cli complete-on-space (off | on)

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set the command-line interface (CLI) to complete a partial command entry when you type a space or a tab. This is the default behavior of the CLI.

Options off—Turn off command completion.

on—Allow either a space or a tab to be used for command completion.

Required Privilege Level view

Related Documentation

- **show cli on page 147**

List of Sample Output **set cli complete-on-space on page 138**

Output Fields When you enter this command, you are provided feedback on the status of your request.

set cli complete-on-space In the following example, pressing the Spacebar changes the partial command entry from **com** to **complete-on-space**. The example shows how adding the keyword **off** at the end of the command disables command completion.

```
user@host> set cli com<Space>
user@host>set cli complete-on-space off
Disabling complete-on-space
```

set cli directory

Syntax set cli directory *directory*

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set the current working directory.

Options *directory*—Pathname of the working directory.

Required Privilege Level view

Related Documentation
- **show cli directory on page 152**

List of Sample Output **set cli directory on page 139**

Output Fields When you enter this command, you are provided feedback on the status of your request.

set cli directory

```
user@host> set cli directory /var/home/regress
Current directory: /var/home/regress
```

set cli idle-timeout

Syntax
set cli idle-timeout
<*minutes*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Set the maximum time that an individual session can be idle before the user is logged off the router or switch.

Options
minutes—(Optional) Maximum idle time. The range of values, in minutes, is 0 through 100,000. If you do not issue this command, and the user's login class does not specify this value, the user is never forced off the system after extended idle times. Setting the value to 0 disables the timeout.

Required Privilege Level
view

Related Documentation
- **show cli on page 147**

List of Sample Output
set cli idle-timeout on page 140

Output Fields
When you enter this command, you are provided feedback on the status of your request.

set cli idle-timeout

```
user@host> set cli idle-timeout 60
Idle timeout set to 60 minutes
```

set cli prompt

Syntax set cli prompt *string*

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set the prompt so that it is displayed within the CLI.

Options *string*—CLI prompt string. To include spaces in the prompt, enclose the string in quotation marks. By default, the string is ***username@hostname***.

Required Privilege Level view

Related Documentation
- **show cli on page 147**

List of Sample Output **set cli prompt on page 141**

Output Fields When you enter this command, the new CLI prompt is displayed.

set cli prompt

```
user@host> set cli prompt lab1-router>
lab1-router>
```

set cli restart-on-upgrade

Syntax set cli restart-on-upgrade string (off | on)

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For an individual session, set the CLI to prompt you to restart the router or switch after upgrading the software.

Options off—Disables the prompt.

on—Enables the prompt.

Required Privilege Level view

Related Documentation
- **show cli on page 147**

List of Sample Output **set cli restart-on-upgrade on page 142**

Output Fields When you enter this command, you are provided feedback on the status of your request.

set cli restart-on-upgrade

```
user@host> set cli restart-on-upgrade on
Enabling restart-on-upgrade
```

set cli screen-length

Syntax set cli screen-length *length*

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set terminal screen length.

Options *length*—Number of lines of text that the terminal screen displays. The range of values, in number of lines, is **24** through **100,000**. The default is **24**.

Additional Information The point at which the **---(more)---** prompt appears on the screen is a function of this setting and the settings for the **set cli screen-width** and **set cli terminal** commands.

Required Privilege Level view

Related Documentation

- **set cli screen-width on page 144**
- **set cli terminal on page 145**
- **show cli on page 147**

List of Sample Output **set cli screen-length on page 143**

Output Fields When you enter this command, you are provided feedback on the status of your request.

set cli screen-length

```
user@host> set cli screen-length 75
Screen length set to 75
```

set cli screen-width

Syntax set cli screen-width *width*

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set the terminal screen width.

Options *width*—Number of characters in a line. The range of values is **80** through **100,000**. The default is **80**.

Additional Information The point at which the **---(more)---** prompt appears on the screen is a function of this setting and the settings for the **set cli screen-length** and **set cli terminal** commands.

Required Privilege Level view

Related Documentation

- **set cli screen-length on page 143**
- **set cli terminal on page 145**
- **show cli on page 147**

List of Sample Output **set cli screen-width on page 144**

Output Fields When you enter this command, you are provided feedback on the status of your request.

set cli screen-width

```
user@host> set cli screen-width
Screen width set to 132
```

set cli terminal

Syntax set cli terminal *terminal-type*

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set the terminal type.

Options *terminal-type*—Type of terminal that is connected to the Ethernet management port:

- **ansi**—ANSI-compatible terminal (80 characters by 24 lines)
- **small-xterm**—Small xterm window (80 characters by 24 lines)
- **vt100**—VT100-compatible terminal (80 characters by 24 lines)
- **xterm**—Large xterm window (80 characters by 65 lines)

Required Privilege Level view

Related Documentation

- **show cli on page 147**

List of Sample Output **set cli terminal on page 145**

Output Fields This command provides no output.

set cli terminal

```
user@host> set cli terminal xterm
```

set cli timestamp

Syntax set cli timestamp (format *timestamp-format* | disable)

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set a timestamp for CLI output.

Options format *timestamp-format*—Set the date and time format for the timestamp. The timestamp format you specify can include the following placeholders in any order:

- **%m**—Two-digit month
- **%d**—Two-digit date
- **%T**—Six-digit hour, minute, and seconds

disable—Remove the timestamp from the CLI.

Required Privilege Level view

Related Documentation

- **show cli on page 147**

List of Sample Output **set cli timestamp on page 146**

Output Fields When you enter this command, you are provided feedback on the status of your request.

set cli timestamp

```
user@host> set cli timestamp format '%m-%d-%T'
'04-21-17:39:13'
CLI timestamp set to: '%m-%d-%T'
```

show cli

Syntax show cli

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display configured CLI settings.

Options This command has no options.

Required Privilege Level view

List of Sample Output **show cli on page 147**

Output Fields Table 27 on page 147 lists the output fields for the **show cli** command. Output fields are listed in the approximate order in which they appear.

Table 27: show cli Output Fields

Field Name	Field Description
CLI complete-on-space	Capability to complete a partial command entry when you type a space or a tab: **on** or **off**.
CLI idle-timeout	Maximum time that an individual session can be idle before the user is logged off the router or switch. When this feature is enabled, the number of minutes is displayed. Otherwise, the state is **disabled**.
CLI restart-on-upgrade	CLI is set to prompt you to restart the router or switch after upgrading the software: **on** or **off**.
CLI screen-length	Number of lines of text that the terminal screen displays.
CLI screen-width	Number of characters in a line on the terminal screen.
CLI terminal	Terminal type.
CLI is operating in	Mode: **enhanced**.
CLI timestamp	Date and time format for the timestamp. If the timestamp is not set, the state is **disabled**.
CLI working directory	Pathname of the working directory.

show cli

```
user@host> show cli
CLI complete-on-space set to on
CLI idle-timeout disabled
CLI restart-on-upgrade set to on
CLI screen-length set to 47
CLI screen-width set to 132
CLI terminal is 'vt100'
CLI is operating in enhanced mode
CLI timestamp disabled
CLI working directory is '/var/home/regress'
```

show cli authorization

Syntax show cli authorization

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display the permissions for the current user.

Options This command has no options.

Required Privilege Level view

List of Sample Output **show cli authorization on page 150**

Output Fields Table 28 on page 149 lists the output fields for the **show cli authorization** command. In the table, all possible permissions are displayed and output fields are listed in alphabetical order.

Table 28: show cli authorization Output Fields

Field Name	Field Description
access	Can view access configuration information.
access-control	Can modify access configuration.
admin	Can view user account information.
admin-control	Can modify user account information.
clear	Can clear learned network information.
configure	Can enter configuration mode.
control	Can modify any configuration.
edit	Can edit configuration files.
field	Reserved for field (debugging) support.
firewall	Can view firewall configuration information.
firewall-control	Can modify firewall configuration information.
floppy	Can read from and write to removable media.
flow-tap	Can view flow-tap configuration information.
flow-tap-control	Can configure flow-tap configuration information.

Table 28: show cli authorization Output Fields *(continued)*

Field Name	Field Description
interface	Can view interface configuration information.
interface-control	Can modify interface configuration information.
maintenance	Can perform system maintenance.
network	Can access the network by entering the **ping**, **ssh**, **telnet**, and **traceroute** commands.
reset	Can reset or restart interfaces and system processes.
rollback	Can rollback to previous configurations.
routing	Can view routing configuration information.
routing-control	Can modify routing configuration information.
secret	Can view passwords and authentication keys in the configuration.
secret-control	Can modify passwords and authentication keys in the configuration.
security	Can view security configuration information.
security-control	Can modify security configuration information.
shell	Can start a local shell.
snmp	Can view SNMP configuration information.
snmp-control	Can modify SNMP configuration information.
system	Can view system configuration information.
system-control	Can modify system configuration information.
trace	Can view trace file settings information.
trace-control	Can modify trace file settings information.
view	Can view current values and statistics.
view-configuration	Can view all configuration information (not including secrets).

show cli authorization

```
user@host> show cli authorization
Current user: 'remote' login: 'user' class ''
Permissions:
    admin       -- Can view user accounts
```

```
    admin-control-- Can modify user accounts
    clear       -- Can clear learned network information
    configure   -- Can enter configuration mode
    control     -- Can modify any configuration
    edit        -- Can edit full files
    field       -- Special for field (debug) support
    floppy      -- Can read and write from the floppy
    interface   -- Can view interface configuration
    interface-control-- Can modify interface configuration
    network     -- Can access the network
    reset       -- Can reset/restart interfaces and daemons
    routing     -- Can view routing configuration
    routing-control-- Can modify routing configuration
    shell       -- Can start a local shell
    snmp        -- Can view SNMP configuration
    snmp-control-- Can modify SNMP configuration
    system      -- Can view system configuration
    system-control-- Can modify system configuration
    trace       -- Can view trace file settings
    trace-control-- Can modify trace file settings
    view        -- Can view current values and statistics
    maintenance -- Can become the super-user
    firewall    -- Can view firewall configuration
    firewall-control-- Can modify firewall configuration
    secret      -- Can view secret configuration
    secret-control-- Can modify secret configuration
    rollback    -- Can rollback to previous configurations
    security    -- Can view security configuration
    security-control-- Can modify security configuration
    access      -- Can view access configuration
    access-control-- Can modify access configuration
    view-configuration-- Can view all configuration (not including secrets)
    flow-tap    -- Can view flow-tap configuration
    flow-tap-control-- Can configure flow-tap service
Individual command authorization:
    Allow regular expression: none
    Deny regular expression: none
    Allow configuration regular expression: none
    Deny configuration regular expression: none
```

show cli directory

Syntax show cli directory

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display the current working directory.

Options This command has no options.

Required Privilege Level view

List of Sample Output **show cli directory on page 152**

Output Fields Table 29 on page 152 lists the output fields for the **show cli directory** command. Output fields are listed in the approximate order in which they appear.

Table 29: show cli directory Output Fields

Field Name	Field Description
Current directory	Pathname of the current working directory.

show cli directory

```
user@host> show cli directory
Current directory: /var/home/regress
```

show cli history

Syntax show cli history
<count>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display a list of previous CLI commands.

Options none—Display all previous CLI commands.

count—(Optional) Maximum number of commands to display.

Required Privilege Level view

List of Sample Output **show cli history on page 153**

Output Fields Table 30 on page 153 lists the output fields for the **show cli history** command. Output fields are listed in the approximate order in which they appear.

Table 30: show cli history Output Fields

Field Name	Field Description
timestamp	Time at which the command was entered.
command-syntax	Command that was entered.

show cli history

```
user@host> show cli history
  11:14:14 -- show arp
  11:22:10 -- show cli authorization
  11:27:12 -- show cli history
```

start shell

Syntax
start shell (csh | sh)
<user *username*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Exit from the CLI environment and create a UNIX-level shell. To return to the CLI, type **exit** from the shell.

Options
csh—Create a UNIX C shell.

sh—Create a UNIX Bourne shell.

user *username*—(Optional) Start the shell as another user.

Additional Information
When you are in the shell, the shell prompt has the following format:

username@hostname%

An example of the prompt is:

root@router%

Required Privilege Level
shell and maintenance

List of Sample Output
start shell csh on page 154

Output Fields
When you enter this command, you are provided feedback on the status of your request.

start shell csh

```
user@host> start shell csh
%

exit
%

username@hostname% start shell sh
%

exit
user@host>
```

PART 5

Junos OS for J-EX Series Switches System Setup

CHAPTER 14

System Setup Overview

Junos OS—Overview

J-EX Series Switch Software Features Overview

The following tables list the J-EX Series Switches software features and the Junos OS release in which they were introduced:

Related Documentation

- J-EX8208 Switch Hardware Overview on page 27
- J-EX8216 Switch Hardware Overview on page 30

Understanding Software Infrastructure and Processes

Each switch runs Junos OS for J-EX Series Switches on its general-purpose processors. The Junos OS includes processes for Internet Protocol (IP) routing and for managing interfaces, networks, and the chassis.

Junos OS runs on the Routing Engine. The Routing Engine kernel coordinates communication among the Junos OS processes and provides a link to the Packet Forwarding Engine.

With the J-Web interface and the command-line interface (CLI) to Junos OS, you configure switching features and routing protocols and set the properties of network interfaces on your switch. After activating a software configuration, use either the J-Web or CLI user interface to monitor the switch, manage operations, and diagnose protocol and network connectivity problems.

- Routing Engine and Packet Forwarding Engine on page 158
- Junos OS Processes on page 158

Routing Engine and Packet Forwarding Engine

A switch has two primary software processing components:

- Packet Forwarding Engine—Processes packets; applies filters, routing policies, and other features; and forwards packets to the next hop along the route to their final destination.
- Routing Engine—Provides three main functions:
  - Creates the packet forwarding switch fabric for the switch, providing route lookup, filtering, and switching on incoming data packets, then directing outbound packets to the appropriate interface for transmission to the network
  - Maintains the routing tables used by the switch and controls the routing protocols that run on the switch.
  - Provides control and monitoring functions for the switch, including controlling power and monitoring system status.

Junos OS Processes

Junos OS running on the Routing Engine and Packet Forwarding Engine consists of multiple processes that are responsible for individual functions.

The separation of functions provides operational stability, because each process accesses its own protected memory space. In addition, because each process is a separate software package, you can selectively upgrade all or part of Junos OS, for added flexibility.

Table 55 on page 393 describes the primary Junos OS processes.

Related Documentation

- For more information about processes, see the *Junos OS Network Operations Guide* at http://www.juniper.net/techpubs/software/junos/.
- For more information about basic system parameters, supported protocols, and software processes, see the *Junos OS System Basics Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.

CHAPTER 15

Initial Configuration

Connecting and Configuring a J-EX Series Switch (CLI Procedure)

There are two ways to connect and configure a J-EX Series switch: one method is through the console using the CLI and the other is using the J-Web interface. This topic describes the CLI procedure.

NOTE: To run the ezsetup script, the switch must have the factory default configuration as the active configuration. If you have configured anything on the switch and want to run ezsetup, revert to the factory default configuration. See "Reverting to the Default Factory Configuration for the J-EX Series Switch" on page 341.

Before you begin connecting and configuring a J-EX Series switch through the console using the CLI:

- Set the following parameter values in the console server or PC:
  - Baud Rate—9600
  - Flow Control—None
  - Data—8
  - Parity—None

- Stop Bits—1
- DCD State—Disregard

To connect and configure the switch from the console:

1. Connect the console port to a laptop or PC using the RJ-45 to DB-9 serial port adapter. The RJ-45 cable and RJ-45 to DB-9 serial port adapter are supplied with the switch.
   - J-EX4200 switch—The console port is located on the rear panel of the switch.
   - J-EX8200 switch—The console port is located on the Switch Fabric and Routing Engine (SRE) module in slot SRE0 in a J-EX8208 switch or on the Routing Engine (RE) module in slot RE0 in a J-EX8216 switch.
2. At the Junos OS shell prompt **root%**, type **ezsetup**.
3. Enter the hostname. This is optional.
4. Enter the root password you plan to use for this device. You are prompted to re-enter the root password.
5. Enter **yes** to enable services like Telnet and SSH. By default, Telnet is not enabled and SSH is enabled.

   

   NOTE: When Telnet is enabled, you will not be able to log in to a J-EX Series switch through Telnet using root credentials. Root login is allowed only for SSH access.

6. Use the Management Options page to select the management scenario:

   

   NOTE: On J-EX8200 switches, only the out-of-band management option is available.

   - **Configure in-band management**. In this scenario you have the following two options:
     - Use the default VLAN.
     - Create a new VLAN—If you select this option, you are prompted to specify the VLAN name, VLAN ID, management IP address, and default gateway. Select the ports that must be part of this VLAN.
   - **Configure out-of-band management**. Specify the IP address and gateway of the management interface. Use this IP address to connect to the switch.
7. Specify the SNMP Read Community, Location, and Contact to configure SNMP parameters. These parameters are optional.
8. Specify the system date and time. Select the time zone from the list. These options are optional.

The configured parameters are displayed. Enter **yes** to commit the configuration.

The configuration is committed as the active configuration for the switch. You can now log in with the CLI or the J-Web interface to continue configuring the switch. If you use the J-Web interface to continue configuring the switch, the Web session is redirected to the new management IP address. If the connection cannot be made, the J-Web interface displays instructions for starting a J-Web session.

Related Documentation

- Connecting and Configuring a J-EX Series Switch (J-Web Procedure) on page 163
- Installing and Connecting a J-EX4200 Switch
- Installing and Connecting a J-EX8208 Switch
- Installing and Connecting a J-EX8216 Switch

Connecting and Configuring a J-EX Series Switch (J-Web Procedure)

There are two ways to connect and configure a J-EX Series switch: one method is through the console using the CLI and the other is using the J-Web interface. This topic describes the J-Web procedure.

To access the J-Web interface, your management device must have the following software installed:

- Operating system: Microsoft Windows XP Service Pack 3
- Browser version: One of the following. Other browsers might work but are not supported by J-Series platforms.
  - Microsoft Internet Explorer version 7.0
  - Mozilla Firefox version 3.0
- Additional requirements:
  - Only English-language browsers are supported.
  - The browser and the network must be able to receive and process HTTP/1.1 gzip compressed data.

NOTE: Before you begin the configuration, enable a DHCP client on the management PC you will connect to the switch so that the switch can obtain an IP address dynamically.

NOTE: Read the following steps before you begin the configuration. You must complete the initial configuration using EZSetup within 10 minutes. The switch exits EZSetup after 10 minutes and reverts to the factory default configuration, and the PC loses connectivity to the switch. The LCD displays a count-down timer when the switch is in initial setup mode.

To connect and configure the switch using the J-Web interface:

1. Transition the switch into initial setup mode by using the **Menu** and **Enter** buttons located to the right of the LCD panel (see Figure 4 on page 164)

Figure 4: LCD Panel in a J-EX4200 or J-EX8200 Switch

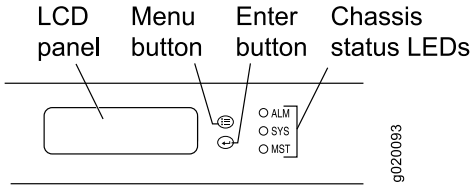

   a. Press the **Menu** button until you see **MAINTENANCE MENU**. Then press the **Enter** button.

   b. Press **Menu** until you see **ENTER EZSetup**. Then press **Enter**.

      If EZSetup does not appear as an option in the menu, select Factory Default to return the switch to the factory default configuration. EZSetup is displayed in the menu only when the switch is set to the factory default configuration.

   c. Press **Enter** to confirm setup and continue with EZSetup.

2. Connect the Ethernet cable from the Ethernet port on the PC to the switch.

   - J-EX4200 switch—Connect the cable to port 0 (**ge-0/0/0**) on the front panel of the switch.
   - J-EX8200 switch—Connect the cable to the port labeled **MGMT** on the Switch Fabric and Routing Engine (SRE) module in slot SRE0 in a J-EX8208 switch or on the Routing Engine (RE) module in slot RE0 in a J-EX8216 switch.

   These ports are configured as the DHCP server with the default IP address, **192.168.1.1**. The switch can assign an IP address to the management PC in the IP address range **192.168.1.2** through **192.168.1.253**.

3. From the PC, open a Web browser, type **http://192.168.1.1** in the address field, and press Enter.

4. On the J-Web login page, type **root** as the username, leave the password field blank, and click **Login**.

5. On the Introduction page, click **Next**.

6. On the Basic Settings page, modify the hostname, the root password, and date and time settings:

   a. Enter the hostname. This is optional.

   b. Enter a password and reenter the password.

   c. Specify the time zone.

d. Synchronize the date and time settings of the switch with the management PC or set them manually by selecting the appropriate option button. This is optional.

e. Click **Next**.

7. Use the Management Options page to select the management scenario:

NOTE: On J-EX8200 switches, only the out-of-band management option is available.

- **In-band Management—Use VLAN 'default' for management**.

  Select this option to configure all data interfaces as members of the default VLAN. Click **Next**. Specify the management IP address and the default gateway for the default VLAN.

- **In-band Management—Create new VLAN for management**.

  Select this option to create a management VLAN. Click **Next**. Specify the VLAN name, VLAN ID, member interfaces, management IP address, and default gateway for the new VLAN.

- **Out-of-band Management—Configure management port**.

  Select this option to configure only the management interface. Click **Next**. Specify the IP address and default gateway for the management interface.

8. Click **Next**.
9. On the Manage Access page, you may select options to enable Telnet, SSH, and SNMP services. For SNMP, you can configure the read community, location, and contact.
10. Click **Next**. The Summary screen displays the configured settings.
11. Click **Finish**.

The configuration is committed as the active switch configuration. You can now log in with the CLI or the J-Web interface to continue configuring the switch.

If you use the J-Web interface to continue configuring the switch, the Web session is redirected to the new management IP address. If the connection cannot be made, the J-Web interface displays instructions for starting a J-Web session.

NOTE: After the configuration is committed, the connectivity between the PC and the switch might be lost. To renew the connection, release and renew the IP address by executing the appropriate commands on the management PC or by removing and reinserting the Ethernet cable.

Related Documentation

- Connecting and Configuring a J-EX Series Switch (CLI Procedure) on page 161
- Installing and Connecting a J-EX4200 Switch

- Installing and Connecting a J-EX8208 Switch
- Installing and Connecting a J-EX8216 Switch

Configuring the LCD Panel on J-EX Series Switches (CLI Procedure)

The LCD panel on the front panel of J-EX Series switches displays a variety of information about the switch in the Status menu and provides the Maintenance menu to allow you to perform basic operations such as initial setup and reboot. You can disable these menus or individual menu options if you do not want switch users to use them. You can also set a custom message that will be displayed on the panel.

This topic describes:

- Disabling or Enabling Menus and Menu Options on the LCD Panel on page 166
- Configuring a Custom Display Message on page 167

Disabling or Enabling Menus and Menu Options on the LCD Panel

By default, the Maintenance menu, the Status menu, and the options in those menus in the LCD panel are enabled. Users can configure and troubleshoot the switch using the Maintenance menu and view certain details about the switch using the Status menu.

If you do not want users to be able to use those menus or use some of the menu options, you can disable the menus or individual menu options. You can re-enable the menus or menu options.

Issue the **show chassis lcd** operational mode command to see which menus and menu options are currently enabled.

NOTE: On some platforms you must specify an FPC slot number in these commands. See the **lcd-menu** statement for details.

To disable a menu:

```
[edit]
user@switch# set chassis lcd-menu menu-item menu-name disable
```

To enable a menu:

```
[edit]
user@switch# delete chassis lcd-menu menu-item menu-name disable
```

To disable a menu option:

```
[edit]
user@switch# set chassis lcd-menu menu-item menu-option disable
```

To enable a menu option:

```
[edit]
user@switch# delete chassis lcd-menu menu-item menu-option disable
```

Configuring a Custom Display Message

You can configure the second line of the LCD to display a custom message temporarily for 5 minutes or permanently.

To display a custom message temporarily:

- On a standalone J-EX4200 switch or a J-EX8200 switch:

  ```
  user@switch> set chassis display message message
  ```

- On a J-EX4200 switch in a Virtual Chassis configuration:

  ```
  user@switch> set chassis display message message fpc-slot slot-number
  ```

To display a custom message permanently:

- On a standalone J-EX4200 switch or a J-EX8200 switch:

  ```
  user@switch> set chassis display message message permanent
  ```

- On a J-EX4200 switch in a Virtual Chassis configuration:

  ```
  user@switch> set chassis display message message fpc-slot slot-number permanent
  ```

NOTE: The Menu button and the Enter button are disabled if the LCD is configured to display a custom message.

To disable the display of the custom message:

```
user@switch> clear chassis display message
```

You can view the custom message by issuing the command **show chassis lcd**.

Related Documentation

- LCD Panel in J-EX4200 Switches
- LCD Panel in a J-EX8200 Switch

Configuring Date and Time for the J-EX Series Switch (J-Web Procedure)

To configure date and time on a J-EX Series switch:

1. Select **Configure** > **System Properties** > **Date & Time**.
2. To modify the information, click **Edit**. Enter information into the Edit Date & Time page as described in Table 31 on page 168.
3. Click one:
   - To apply the configuration, click **OK**.
   - To cancel your entries and return to the System Properties page, click **Cancel**.

NOTE: After you make changes to the configuration in this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options** > **Commit**. See "Using the Commit Options to Commit Configuration Changes (J-Web Procedure)" on page 334 for details about all commit options.

Table 31: Date and Time Settings

Time	Function	Your Action
Time Zone	Identifies the timezone that the switching platform is located in.	Select the appropriate time zone from the list.
Set Time	Synchronizes the system time with that of the NTP server. You can also manually set the system time and date.	To immediately set the time, click one: • **Synchronize with PC time**—The switch synchronizes the time with that of the PC. • **NTP Servers**—The switch sends a request to the NTP server and synchronizes the system time. • **Manual**—A pop-up window allows you to select the current date and time from a list.

Related Documentation

- J-Web User Interface for J-EX Series Switches Overview on page 129

Configuring System Identity for a J-EX Series Switch (J-Web Procedure)

To configure identification details for a J-EX Series switch:

1. Select **Configure** > **System Properties** > **System Identity**. The System Identity page displays configuration details.
2. To modify the configuration, click **Edit**. Enter information into the System Identity page as described in Table 32 on page 169.

NOTE: After you make changes to the configuration in this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options** > **Commit**. See "Using the Commit Options to Commit Configuration Changes (J-Web Procedure)" on page 334 for details about all commit options.

Table 32: Set Up Configuration Summary

Field	Function	Your Action
Host Name	Defines the hostname of the switching platform.	Type the hostname.
Domain Name	Defines the network or subnetwork that the machine belongs to.	Type the domain name.
Root Password	Sets the root password that user root can use to log in to the switching platform.	Type a plain-text password. The system encrypts the password. NOTE: After a root password has been defined, it is required when you log in to the J-Web user interface or the CLI.
Confirm Root Password	Verifies that the root password has been typed correctly.	Retype the password.
DNS Name Servers	Specifies a DNS server for the switching platform to use to resolve hostnames into addresses.	To add an IP address, click **Add**. To edit an IP address, click **Edit**. To delete an IP address, click **Delete**.
Domain Search	Specifies the domains to be searched.	To add a domain, click **Add**. To edit a domain click **Edit**. To delete a domain, click **Delete**.

Related Documentation

- Configuring Date and Time for the J-EX Series Switch (J-Web Procedure) on page 167

CHAPTER 16

Configuration Statements for System Setup

arp

Syntax

```
arp {
    aging-timer minutes;
    passive-learning;
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify ARP options. You can enable backup VRRP routers to learn ARP requests for VRRP-IP to VRRP-MAC address translation. You can also set the time interval between ARP updates.

Options

aging-timer—Time interval in minutes between ARP updates. In environments where the number of ARP entries to update is high (for example, on routers only, metro Ethernet environments), increasing the time between updates can improve system performance.

passive-learning—Configures backup VRRP routers or switches to learn the ARP mappings (IP-to-MAC address) for hosts sending the requests. By default, the backup VRRP router drops these requests; therefore, if the master router fails, the backup router must learn all entries present in the ARP cache of the master router. Configuring passive learning reduces transition delay when the backup router is activated.

Default: 20 minutes

Range: 5 to 240 minutes

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring the Junos OS ARP Learning and Aging Options for Mapping IPv4 Network Addresses to MAC Addresses
- *Junos OS Network Interfaces Configuration Guide*

authentication-key

Syntax authentication-key *key-number* type *type* value *password*;

Hierarchy Level [edit system ntp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure Network Time Protocol (NTP) authentication keys so that the router or switch can send authenticated packets. If you configure the router or switch to operate in authenticated mode, you must configure a key.

Both the keys and the authentication scheme (MD5) must be identical between a set of peers sharing the same key number.

Options ***key-number***—Positive integer that identifies the key.

type *type*—Authentication type. It can only be **md5**.

value *password*—The key itself, which can be from 1 through 8 ASCII characters. If the key contains spaces, enclose it in quotation marks.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring NTP Authentication Keys
- **broadcast on page 174**
- **peer on page 194**
- **server on page 197**
- **trusted-key on page 200**

auxiliary

Syntax
```
auxiliary {
  type terminal-type;
}
```

Hierarchy Level [edit system **ports**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the characteristics of the auxiliary port.

Default The auxiliary port is disabled.

Options **type *terminal-type***—Type of terminal that is connected to the port.

Range: **ansi**, **vt100**, **small-xterm**, **xterm**

Default: The terminal type is unknown, and the user is prompted for the terminal type.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring the Junos OS to Set Console and Auxiliary Port Properties

boot-server (NTP)

Syntax boot-server *address*;

Hierarchy Level [edit system **ntp**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the server that NTP queries when the router or switch boots to determine the local date and time.

When you boot the router or switch, it issues an **ntpdate** request, which polls a network server to determine the local date and time. You need to configure a server that the router or switch uses to determine the time when the router or switch boots. Otherwise, NTP will not be able to synchronize to a time server if the server's time appears to be very far off of the local router's or switch's time.

Options ***address***—Address of an NTP server. You must specify an address, not a hostname.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Synchronizing and Coordinating Time Distribution Using NTP

broadcast

Syntax broadcast *address* <key *key-number*> <version *value*> <ttl *value*>;

Hierarchy Level [edit system **ntp**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the local router or switch to operate in broadcast mode with the remote system at the specified ***address***. In this mode, the local router or switch sends periodic broadcast messages to a client population at the specified broadcast or multicast ***address***. Normally, you include this statement only when the local router or switch is operating as a transmitter.

Options ***address***—The broadcast address on one of the local networks or a multicast address assigned to NTP. You must specify an address, not a hostname. If the multicast address is used, it must be **224.0.1.1**.

key *key-number*—(Optional) All packets sent to the address include authentication fields that are encrypted using the specified key number.

Range: Any unsigned 32-bit integer

ttl *value*—(Optional) Time-to-live (TTL) value to use.

Range: 1 through 255

Default: 1

version *value*—(Optional) Specify the version number to be used in outgoing NTP packets.

Range: 1 through 4

Default: 4

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring the NTP Time Server and Time Services

broadcast-client

Syntax broadcast-client;

Hierarchy Level [edit system ntp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the local router or switch to listen for broadcast messages on the local network to discover other servers on the same subnet.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring the Router or Switch to Listen for Broadcast Messages Using NTP

console (Physical Port)

Syntax

```
console {
  disable;
  insecure;
  log-out-on-disconnect;
  type terminal-type;
}
```

Hierarchy Level [edit system **ports**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the characteristics of the console port.

Default The console port is enabled and its speed is 9600 baud.

Options

disable—Disable console login connections.

insecure—Disable root login connections to the console and auxiliary ports. Configuring the console port as insecure also prevents superusers and anyone with a user identifier (UID) of 0 from establishing terminal connections in multiuser mode.

log-out-on-disconnect—Log out the session when the data carrier on the console port is lost.

type *terminal-type*—Type of terminal that is connected to the port.

Range: **ansi**, **vt100**, **small-xterm**, **xterm**

Default: The terminal type is unknown, and the user is prompted for the terminal type.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring the Junos OS to Set Console and Auxiliary Port Properties

default-address-selection

Syntax default-address-selection;

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Use the loopback interface, **lo0**, as the source address for all locally generated IP packets. The **lo0** interface is the interface to the router's or switch's Routing Engine.

Default The outgoing interface is used as the source address.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring the Junos OS to Select a Fixed Source Address for Locally Generated TCP/IP Packets
- *Junos OS Network Interfaces Configuration Guide*

domain-name (Router)

Syntax domain-name *domain-name*;

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the name of the domain in which the router or switch is located. This is the default domain name that is appended to hostnames that are not fully qualified.

Options ***domain-name***—Name of the domain.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring the Domain Name for the Router or Switch

gre-path-mtu-discovery

Syntax (gre-path-mtu-discovery | no-gre-path-mtu-discovery);

Hierarchy Level [edit system internet-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure path MTU discovery for outgoing GRE tunnel connections:

- **gre-path-mtu-discovery**—Path MTU discovery is enabled.
- **no-gre-path-mtu-discovery**—Path MTU discovery is disabled.

Default Path MTU discovery is enabled.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring the Junos OS for Path MTU Discovery on Outgoing GRE Tunnel Connections

host-name

Syntax host-name *hostname*;

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set the hostname of the router or switch.

Options ***hostname***—Name of the router or switch.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring the Hostname of the Router or Switch

icmpv4-rate-limit

Syntax

```
icmpv4-rate-limit {
    bucket-size seconds;
    packet-rate pps;
}
```

Hierarchy Level [edit system **internet-options**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure rate-limiting parameters for ICMPv4 messages sent.

Options

bucket-size *seconds*—Number of seconds in the rate-limiting bucket.

Range: 0 through 4294967295 seconds

Default: 5

packet-rate *pps*—Rate-limiting packets earned per second.

Range: 0 through 4294967295 pps

Default: 1000

Required Privilege Level
admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring the Junos OS ICMPv4 Rate Limit for ICMPv4 Routing Engine Messages

icmpv6-rate-limit

Syntax

```
icmpv6-rate-limit {
    bucket-size seconds;
    packet-rate packet-rate;
}
```

Hierarchy Level [edit system internet-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure rate-limiting parameters for ICMPv6 messages sent.

Options

bucket-size *seconds*—Number of seconds in the rate-limiting bucket.

Range: 0 through 4294967295 seconds

Default: 5

packet-rate *pps*—Rate-limiting packets earned per second.

Range: 0 through 4294967295 pps

Default: 1000

Required Privilege Level
admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring the Junos OS ICMPv6 Rate Limit for ICMPv6 Routing Engine Messages

inet6-backup-router

Syntax inet6-backup-router *address* <destination *destination-address*>;

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set a default router (running IP version 6 [IPv6]) to use while the local router or switch (running IPv6) is booting and if the routing protocol processes fail to start. The Junos OS removes the route to this router or switch as soon as the software starts.

Options ***address***—Address of the default router.

destination *destination-address*—(Optional) Destination address that is reachable through the backup router. Include this option to achieve network reachability while loading, configuring, and recovering the router or switch, but without the risk of installing a default route in the forwarding table.

Default: All hosts (default route) are reachable through the backup router.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring a Backup Router

internet-options

Syntax

```
internet-options {
  (gre-path-mtu-discovery | no-gre-path-mtu-discovery);
  icmpv4-rate-limit bucket-size bucket-size packet-rate packet-rate;
  icmpv6-rate-limit bucket-size bucket-size packet-rate packet-rate;
  (ipip-path-mtu-discovery | no-ipip-path-mtu-discovery);
  ipv6-duplicate-addr-detection-transmits;
  (ipv6-reject-zero-hop-limit | no-ipv6-reject-zero-hop-limit);
  (ipv6-path-mtu-discovery | no-ipv6-path-mtu-discovery);
  ipv6-path-mtu-discovery-timeout;
  no-tcp-rfc1323;
  no-tcp-rfc1323-paws;
  (path-mtu-discovery | no-path-mtu-discovery);
  source-port upper-limit <upper-limit>;
  (source-quench | no-source-quench);
  tcp-drop-synfin-set;
  tcp-mss mss-value;
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure system IP options to protect against certain types of DoS attacks.

The remaining statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring the Junos OS ICMPv4 Rate Limit for ICMPv4 Routing Engine Messages
- Configuring the Junos OS ICMPv6 Rate Limit for ICMPv6 Routing Engine Messages
- Configuring the Junos OS for IP-IP Path MTU Discovery on IP-IP Tunnel Connections
- Configuring the Junos OS for Path MTU Discovery on Outgoing GRE Tunnel Connections
- Configuring the Junos OS for Path MTU Discovery on Outgoing TCP Connections
- Configuring the Junos OS for IPv6 Duplicate Address Detection Attempts
- Configuring the Junos OS for Acceptance of IPv6 Packets with a Zero Hop Limit
- Configuring the Junos OS to Ignore ICMP Source Quench Messages
- Configuring the Junos OS to Enable the Router or Switch to Drop Packets with the SYN and FIN Bits Set
- Configuring the Junos OS to Disable TCP RFC 1323 Extensions
- Configuring the Junos OS to Disable the TCP RFC 1323 PAWS Extension
- Configuring the Junos OS to Extend the Default Port Address Range
- Configuring TCP MSS for Session Negotiation

ipip-path-mtu-discovery

Syntax (ipip-path-mtu-discovery | no-ipip-path-mtu-discovery);

Hierarchy Level [edit system internet-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure path MTU discovery for outgoing IP-IP tunnel connections:

- **ipip-path-mtu-discovery**—Path MTU discovery is enabled.
- **no-ipip-path-mtu-discovery**—Path MTU discovery is disabled.

Default Path MTU discovery is enabled.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring the Junos OS for IP-IP Path MTU Discovery on IP-IP Tunnel Connections

ipv6-duplicate-addr-detection-transmits

Syntax ipv6-duplicate-addr-detection-transmits;

Hierarchy Level [edit system internet-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Control the number of attempts for IPv6 duplicate address detection.

Default The default value is 3.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring the Junos OS for IPv6 Duplicate Address Detection Attempts

ipv6-path-mtu-discovery

Syntax (ipv6-path-mtu-discovery | no-ipv6-path-mtu-discovery);

Hierarchy Level [edit system internet-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure path MTU discovery for IPv6 packets:

- **ipv6-path-mtu-discovery**—IPv6 path MTU discovery is enabled.
- **no-ipv6-path-mtu-discovery**—IPv6 path MTU discovery is disabled.

Default IPv6 path MTU discovery is enabled.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring the Junos OS for IPv6 Path MTU Discovery

ipv6-path-mtu-discovery-timeout

Syntax ipv6-path-mtu-discovery-timeout *minutes*;

Hierarchy Level [edit system internet-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set the IPv6 path MTU discovery timeout interval.

Options ***minutes***—IPv6 path MTU discovery timeout.
Default: 10 minutes

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring the Junos OS for IPv6 Path MTU Discovery

ipv6-reject-zero-hop-limit

Syntax (ipv6-reject-zero-hop-limit | no-ipv6-reject-zero-hop-limit);

Hierarchy Level [edit system internet-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable and disable rejecting incoming IPv6 packets with a zero hop limit value in their header.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring the Junos OS for Acceptance of IPv6 Packets with a Zero Hop Limit

lcd-menu

Syntax

J-EX4200 switch:

```
lcd-menu fpc slot-number {
  menu-item (menu-name | menu-option);
}
```

J-EX8200 switch:

```
lcd-menu {
  menu-item (menu-name | menu-option);
}
```

Hierarchy Level [edit chassis]

Release Information Statement introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Disable or enable the Maintenance menu or the Status menu in the LCD panel.

Options

none—(J-EX8200 switches only) Disable or enable the specified menu or menu options.

fpc *slot-number*—(J-EX4200 switches only) Disable or enable the specified menu or menu options, where ***slot-number*** is:

- 0—On a standalone J-EX4200 switch
- 0–9—On a J-EX4200 switch in a Virtual Chassis. The value is the member ID of the switch.

The remaining statement is explained separately.

Required Privilege Level

view-level—To view this statement in the configuration.
control-level—To add this statement to the configuration.

Related Documentation

- Configuring the LCD Panel on J-EX Series Switches (CLI Procedure) on page 166
- LCD Panel in J-EX4200 Switches
- LCD Panel in a J-EX8200 Switch

location

Syntax

```
location {
  altitude feet;
  building name;
  country-code code;
  floor number;
  hcoord horizontal-coordinate;
  lata service-area;
  latitude degrees;
  longitude degrees;
  npa-nxx number;
  postal-code postal-code;
  rack number;
  vcoord vertical-coordinate;
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the system location in various formats.

Options

altitude *feet*—Number of feet above sea level.

building *name*—Name of building. The name of the building can be 1 to 28 characters in length. If the string contains spaces, enclose it in quotation marks (" ").

country-code *code*—Two-letter country code.

floor *number*—Floor in the building.

hcoord *horizontal-coordinate*—Bellcore Horizontal Coordinate.

lata *service-area*—Long-distance service area.

latitude *degrees*—Latitude in degree format.

longitude *degrees*—Longitude in degree format.

npa-nxx *number*—First six digits of the phone number (area code and exchange).

postal-code *postal-code*—Postal code.

rack *number*—Rack number.

vcoord *vertical-coordinate*—Bellcore Vertical Coordinate.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring the Physical Location of the Router or Switch

menu-item

Syntax menu-item (*menu-name* | *menu-option*);

Hierarchy Level [edit chassis **lcd-menu**]

Release Information Statement introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Disable or enable the Maintenance menu, the Status menu, or an individual option in one of those menus in the LCD panel.

Options ***menu-name***—Name of the LCD menu:

- **maintenance-menu**
- **status-menu**

menu-option—Specific option on one of the LCD menus. You must include the quotation marks when you type the option.

- **"maintenance-menu halt-menu"**—System halt option
- **"maintenance-menu system-reboot"**—System reboot option
- **"maintenance-menu rescue-config"**—Load rescue option
- **"maintenance-menu vc-uplink-config"**—(J-EX4200 switches only) Request VC port option for a J-EX4200 switch in a Virtual Chassis configuration
- **"maintenance-menu factory-default"**—Factory default option
- **"status-menu vcp-status"**—(J-EX4200 switches only) Virtual Chassis port (VCP) status for a J-EX4200 switch in a Virtual Chassis configuration
- **"status-menu sf-status1-menu"**—(J-EX8200 switches only) Status of the switch fabric on the Switch Fabric and Routing Engine (SRE) module in slot SRE0 on J-EX8208 switches. Status of the switch fabric on the Switch Fabric (SF) modules in slots SF0 and SF1 on J-EX8216 switches.
- **"status-menu sf-status2-menu"**—(J-EX8200 switches only) Status of the switch fabric on the SRE module in slot SRE1 on J-EX8208 switches. Status of the switch fabric on the SF modules in slots SF2–SF5 on J-EX8216 switches.
- **"status-menu sf-status3-menu"**—(J-EX8216 switches only) Status of the switch fabric on the SF modules in slots SF6 and SF7
- **"status-menu power-status"**—(J-EX4200 switches only) Status of the power supply
- **"status-menu psu-status1-menu"**—(J-EX8200 switches only) Status of the power supplies in slots P0 and P1
- **"status-menu psu-status2-menu"**—(J-EX8200 switches only) Status of the power supplies in slots P2–P5
- **"status-menu environ-status"**—Status of the fan and the temperature

- **"status-menu show-version"**—The version of Junos OS for J-EX Series switches loaded on the switch

Required Privilege Level
view-level—To view this statement in the configuration.
control-level—To add this statement to the configuration.

Related Documentation
- Configuring the LCD Panel on J-EX Series Switches (CLI Procedure) on page 166
- LCD Panel in J-EX4200 Switches
- LCD Panel in a J-EX8200 Switch

multicast-client

Syntax
multicast-client <*address*>;

Hierarchy Level
[edit system **ntp**]

Release Information
Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
For NTP, configure the local router or switch to listen for multicast messages on the local network to discover other servers on the same subnet.

Options
address—(Optional) One or more IP addresses. If you specify addresses, the router or switch joins those multicast groups.
Default: 224.0.1.1.

Required Privilege Level
system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring the Router or Switch to Listen for Multicast Messages Using NTP

no-multicast-echo

Syntax no-multicast-echo;

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable the Routing Engine from responding to ICMP echo requests sent to multicast group addresses.

Default The Routing Engine responds to ICMP echo requests sent to multicast group addresses.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring the Junos OS to Disable the Routing Engine Response to Multicast Ping Packets

no-ping-record-route

Syntax no-ping-record-route;

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the Junos OS to disable the reporting of the IP address in ping responses.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring the Junos OS to Disable the Reporting of IP Address and Timestamps in Ping Responses

no-ping-time-stamp

Syntax no-ping-time-stamp;

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the Junos OS to disable the recording of timestamps in ping responses.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring the Junos OS to Disable the Reporting of IP Address and Timestamps in Ping Responses

no-redirects

Syntax no-redirects;

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable the sending of protocol redirect messages by the router or switch.

To disable the sending of redirect messages on a per-interface basis, include the **no-redirects** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family *family*]** hierarchy level.

Default The router or switch sends redirect messages.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring the Junos OS to Disable Protocol Redirect Messages on the Router or Switch
- *Junos OS Network Interfaces Configuration Guide*

no-tcp-rfc1323

Syntax no-tcp-rfc1323;

Hierarchy Level [edit system internet-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the Junos OS to disable RFC 1323 TCP extensions.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring the Junos OS to Disable TCP RFC 1323 Extensions

no-tcp-rfc1323-paws

Syntax no-tcp-rfc1323-paws;

Hierarchy Level [edit system internet-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the Junos OS to disable the RFC 1323 Protection Against Wrapped Sequence (PAWS) number extension.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring the Junos OS to Disable the TCP RFC 1323 PAWS Extension

ntp

Syntax

```
ntp {
  authentication-key number type type value password;
  boot-server address;
  broadcast <address> <key key-number> <version value> <ttl value>;
  broadcast-client;
  multicast-client <address>;
  peer address <key key-number> <version value> <prefer>;
  server address <key key-number> <version value> <prefer>;
  source-address source-address;
  trusted-key [ key-numbers ];
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure NTP on the router or switch.

The remaining statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Synchronizing and Coordinating Time Distribution Using NTP

path-mtu-discovery

Syntax (path-mtu-discovery | no-path-mtu-discovery);

Hierarchy Level [edit system internet-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure path MTU discovery for outgoing Transmission Control Protocol (TCP) connections:

- **path-mtu-discovery**—Path MTU discovery is enabled.
- **no-path-mtu-discovery**—Path MTU discovery is disabled.

Default Path MTU discovery is enabled.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring the Junos OS for Path MTU Discovery on Outgoing TCP Connections

peer

Syntax peer *address* <key *key-number*> <version *value*> <prefer>;

Hierarchy Level [edit system **ntp**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For NTP, configure the local router or switch to operate in symmetric active mode with the remote system at the specified address. In this mode, the local router or switch and the remote system can synchronize with each other. This configuration is useful in a network in which either the local router or switch or the remote system might be a better source of time.

Options ***address***—Address of the remote system. You must specify an address, not a hostname.

key *key-number*—(Optional) All packets sent to the address include authentication fields that are encrypted using the specified key number.

Range: Any unsigned 32-bit integer

prefer—(Optional) Mark the remote system as the preferred host, which means that if all other factors are equal, this remote system is chosen for synchronization among a set of correctly operating systems.

version *value*—(Optional) Specify the NTP version number to be used in outgoing NTP packets.

Range: 1 through 4

Default: 4

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring the NTP Time Server and Time Services

ports

Syntax

```
ports {
  auxiliary {
    type terminal-type;
  }
  console {
    type terminal-type;
  }
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the properties of the console and auxiliary ports. The ports are located on the router's craft interface.

See the switch's hardware documentation for port locations.

The remaining statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring the Junos OS to Set Console and Auxiliary Port Properties

processes

Syntax

```
processes {
  process-name (enable | disable) failover (alternate-media | other-routing-engine);
  timeout seconds;
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure which Junos OS processes are running on the router or switch.

CAUTION: Never disable any of the software processes unless instructed to do so by a customer support engineer.

Default All processes are enabled by default.

Options

(enable | disable)—(Optional) Enable or disable a specified process.

failover (alternate-media | other-routing-engine)—(Optional) For routers or switches with redundant Routing Engines only, switch to backup media if a process fails repeatedly. If a process fails four times within 30 seconds, the router or switch reboots from the alternate media or the other Routing Engine.

process-name—One of the valid process names. You can obtain a complete list of process names by using the CLI command completion feature. After specifying a process name, command completion also indicates any additional options for that process.

timeout *seconds*—(Optional) How often the system checks the watchdog timer, in seconds. If the watchdog timer has not been checked in the specified number of seconds, the system reloads. If you set the time value too low, it is possible for the system to reboot immediately after it loads.

Values: **15**, **60**, or **180**

Default: **180** seconds (rounded up to 291 seconds by the Junos OS kernel)

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Disabling Junos OS Processes

server (NTP)

Syntax server *address* <key *key-number*> <version *value*> <prefer>;

Hierarchy Level [edit system ntp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For NTP, configure the local router or switch to operate in client mode with the remote system at the specified ***address***. In this mode, the local router or switch can be synchronized with the remote system, but the remote system can never be synchronized with the local router or switch.

Options ***address***—Address of the remote system. You must specify an address, not a hostname.

key ***key-number***—(Optional) Use the specified key number to encrypt authentication fields in all packets sent to the specified address.

Range: Any unsigned 32-bit integer

prefer—(Optional) Mark the remote system as preferred host, which means that if all other things are equal, this remote system is chosen for synchronization among a set of correctly operating systems.

version *value*—(Optional) Specify the version number to be used in outgoing NTP packets.

Range: 1 through 4

Default: 4

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring the NTP Time Server and Time Services

tcp-drop-synfin-set

Syntax tcp-drop-synfin-set;

Hierarchy Level [edit system internet-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the router or switch to drop packets that have both the SYN and FIN bits set.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring the Junos OS to Enable the Router or Switch to Drop Packets with the SYN and FIN Bits Set

traceoptions (SBC Configuration Process)

Syntax

```
traceoptions {
  file filename <files number> <match regex> <size size>
    <world-readable | no-world-readable>;
  flag flag;
}
```

Hierarchy Level [edit system processes sbc-configuration-process]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure trace options for the session border controller (SBC) process of the border signaling gateway (BSG).

Options

file *filename*—Name of the file that receives the output of the tracing operation. Enclose the name in quotation marks. All files are placed in the directory **/var/log**. You can include the following file options:

- **files *number***—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

  If you specify a maximum number of files, you must also specify a maximum file size with the **size** option and a filename.

  Range: 2 through 1000
  Default: 3 files

- **match *regex***—(Optional) Refine the output to include lines that contain the regular expression.
- **no-world-readable**—(Optional) Disable unrestricted file access.
- **size *size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the trace-file again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum file size, you also must specify a maximum number of trace files with the files option and filename.

  Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB.
  Range: 10 KB through 1 GB
  Default: 128 KB

- **world-readable**—(Optional) Enable unrestricted file access.

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **all** ***trace-level***—Trace all SBC process operations.
- **common** ***trace-level***—Trace common events.
- **configuration** ***trace-level***—Trace configuration events.
- **device-monitor** ***trace-level***—Trace device monitor events.
- **ipc** ***trace-level***—Trace IPC events.
- **memory—pool** ***trace-level***—Trace memory pool events.
- ***trace-level***—Trace level options are related to the severity of the event being traced. When you choose a trace level, messages at that level and higher levels are captured. Enter one of the following trace levels as the ***trace-level***:
  - **debug**—Log all code flow of control.
  - **error**—Log failures with a short-term effect.
  - **info**—Log summary for normal operations, such as the policy decisions made for a call.
  - **trace**—Log program trace START and EXIT macros.
  - **warning**—Log failure recovery events or failure of an external entity.
- **ui** ***trace-level***—Trace user interface operations.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- See “Troubleshooting the IMSG” in the *Junos OS Multiplay Solutions Guide*
- System Management Configuration Statements

trusted-key

Syntax trusted-key [*key-numbers*];

Hierarchy Level [edit system ntp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For NTP, configure the keys you are allowed to use when you configure the local router or switch to synchronize its time with other systems on the network.

Options ***key-numbers***—One or more key numbers. Each key can be any 32-bit unsigned integer except 0.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring NTP Authentication Keys
- **authentication-key on page 172**
- **broadcast on page 174**
- **peer on page 194**
- **server on page 197**

CHAPTER 17

Operational Mode Commands for System Setup

clear chassis display message

Syntax clear chassis display message

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Clear or stop a text message on the craft interface display, which is on the front of the router or on the LCD panel display on the switch. The craft interface alternates the display of text messages with standard craft interface messages, switching between messages every 2 seconds. By default, on both the router and the switch, the text message is displayed for 5 minutes. The craft interface display has four 20-character lines. The LCD panel display has two 16-character lines, and text messages appear only on the second line.

Options none—Clear or stop a text message on the craft interface display.

Required Privilege Level clear

Related Documentation

- Configuring the LCD Panel Display on J-EX Series Switches (CLI Procedure) on page 166
- **set chassis display message on page 233**
- show chassis craft-interface

List of Sample Output **clear chassis display message on page 202**

Output Fields See show chassis craft-interface for an explanation of output fields.

clear chassis display message The following example displays and then clears the text message on the craft interface display:

```
user@host> show chassis craft-interface
Red alarm:     LED off, relay off
Yellow alarm:  LED off, relay off
Host OK LED:   On
Host fail LED: Off
FPCs     0  1  2  3  4  5  6  7
-------------------------------
Green  ..  *..  *  *.
Red    ........
LCD screen:
     +--------------------+
     |NOC contact Dusty   |
     |(888) 526-1234      |
     +--------------------+

user@host> clear chassis display message

user@host> show chassis craft-interface
Red alarm:     LED off, relay off
Yellow alarm:  LED off, relay off
Host OK LED:   On
Host fail LED: Off
FPCs     0  1  2  3  4  5  6  7
```

```
-------------------------------
Green  ..  *..  *  *.
Red    ........
LCD screen:
     +--------------------+
     |host                |
     |Up: 0+17:05:47      |
     |                    |
     |Temperature OK      |
     +--------------------+
```

clear system reboot

Syntax

clear system reboot
<both-routing-engines>

Syntax (J-EX Series Switch)

clear system reboot
<all-members>
<both-routing-engines>
<local>
<member *member-id*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Clear any pending system software reboots or halts.

Options

none—Clear all pending system software reboots or halts.

all-members—(J-EX4200 switches only) (Optional) Clear all halt or reboot requests on all members of the Virtual Chassis configuration.

both-routing-engines—(Systems with multiple Routing Engines) (Optional) Clear all halt or reboot requests on both Routing Engines. On a TX Matrix router, clear both Routing Engines on all chassis connected to the TX Matrix router. Likewise, on a TX Matrix Plus router, clear both Routing Engines on all chassis connected to the TX Matrix Plus router.

local—(J-EX4200 switches only) (Optional) Clear all halt or reboot requests on the local Virtual Chassis member.

member *member-id*—(J-EX4200 switches only) (Optional) Clear all halt or reboot requests on the specified member of the Virtual Chassis configuration. Replace ***member-id*** with a value from 0 through 9.

Required Privilege Level

maintenance

Related Documentation

- **request system reboot on page 101**

List of Sample Output

clear system reboot on page 205

Output Fields

When you enter this command, you are provided feedback on the status of your request.

clear system reboot

```
user@host> clear system reboot
reboot requested by root at Sat Dec 12 19:37:34 1998
[process id 17855]
Terminating...
```

configure

Syntax

configure
<dynamic>
<exclusive>
<private>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Enter configuration mode. When this command is entered without any optional keywords, everyone can make configuration changes and commit all changes made to the configuration.

Options

none—Enter configuration mode.

dynamic—(Optional) Configure routing policies and certain routing policy objects in a dynamic database that is not subject to the same verification required in the standard configuration database. As a result, the time it takes to commit changes to the dynamic database is much shorter than for the standard configuration database. You can then reference these policies and policy objects in routing policies you configure in the standard database.

exclusive—(Optional) Lock the candidate configuration for as long as you remain in configuration mode, allowing you to make changes without interference from other users. Other users can enter and exit configuration mode, but they cannot change the configuration.

private—(Optional) Allow multiple users to edit different parts of the configuration at the same time and to commit only their own changes, or to roll back without interfering with one another's changes. You cannot commit changes in configure private mode when another user is in configure exclusive mode.

Additional Information

For more information about the different methods of entering configuration mode and the restrictions that apply, see the *Junos OS System Basics Configuration Guide*.

Required Privilege Level

configure

Related Documentation

- **show configuration on page 244**

List of Sample Output

configure on page 206

Output Fields

When you enter this command, you are placed in configuration mode and the system prompt changes from ***hostname>*** to ***hostname#***.

configure

```
user@host> configure
Entering configuration mode
[edit]
user@host#
```

op

Syntax

op *filename*
<detail>
<*argument-name argument-value*>
<key (md5 | sha-256 | sha1) *key-value*
<url *url*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Execute an op script stored in one of the following locations:

- On the router or switch in the **/var/db/scripts/op** directory
- At a remote URL

Options

detail—(Optional) Display detailed output.

argument-name argument-value—(Optional) Specify one or more arguments to the script. For each argument you include on the command line, you must specify a corresponding value for the argument.

key (md5 | sha-256 | sha1) *key-value*—(Optional) With the <url> option, specify a checksum hash to verify the integrity of the script. You can include the <key> option if the **checksum** statement is included at the **[edit system scripts op file *filename*]** hierarchy level.

url *url*—(Optional) Specify a URL where the script is located.

Additional Information

For more information about Junos OS op scripts, see the *Junos OS Configuration and Operations Automation Guide*.

Required Privilege Level

maintenance

Related Documentation

- Executing an Op Script in the *Junos OS Configuration and Operations Automation Guide*
- Executing an Op Script from a Remote Site in the *Junos OS Configuration and Operations Automation Guide*
- checksum
- **file checksum md5 on page 364**
- **file checksum sha-256 on page 366**
- **file checksum sha1 on page 365**

List of Sample Output

op on page 208
op url on page 208

Output Fields When you enter this command, you are provided feedback on the status of your request.

op `user@host>` **op script1 interface ge-0/2/0.0 protocol inet**

op url `user@host>` **op url https://www.juniper.net/fa/2009-04-01.01.slax key md5 8de24d09e1d90b2581bb937d2a5ad590 interface ge-0/2/0.0 protocol inet**

request chassis pic

Syntax request chassis pic (offline | online) fpc-slot *slot-number* pic-slot *slot-number*

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Control the operation of the PIC.

NOTE: To view a list of built-in PICs on the router or switch chassis, use the show chassis hardware command.

Options offline—Take the PIC offline.

online—Bring the PIC online.

fpc-slot *slot-number*—Flexible PIC Concentrator (FPC) slot number. Replace ***slot-number*** with a value appropriate for your router or switch:

- J-EX Series switches:
  - J-EX4200 standalone switches—0.
  - J-EX4200 switches in a Virtual Chassis configuration—0 through 9 (switch's member ID).
  - J-EX8208 switches—0 through 7 (line card).
  - J-EX8216 switches—0 through 15 (line card).

pic-slot *slot-number*—PIC slot number. For J-EX4200 switches, it is 0 for built-in network interfaces and 1 for interfaces on uplink modules. For J-EX8208 and J-EX8216 switches, it is 0.

Required Privilege Level maintenance

Related Documentation

- **show chassis hardware on page 643**
- **show chassis pic on page 650**

List of Sample Output **request chassis pic on page 209**

Output Fields When you enter this command, you are provided feedback on the status of your request.

request chassis pic

```
user@host> request chassis pic pic-slot 0 online fpc-slot 0
FPC 0, PIC 0 is already online
```

request chassis routing-engine master

Syntax request chassis routing-engine master (acquire | release | switch)
<force>
<no-confirm>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For routers or switches with multiple Routing Engines, control which Routing Engine is the master.

NOTE: Successive graceful Routing Engine switchover events must be a minimum of 240 seconds (4 minutes) apart after both Routing Engines have come up.

If the router or switch displays a warning message similar to "Standby Routing Engine is not ready for graceful switchover. Packet Forwarding Engines that are not ready for graceful switchover might be reset," do not attempt switchover. If you choose to proceed with switchover, only the Packet Forwarding Engines that were not ready for graceful switchover are reset. None of the Flexible PIC concentrators (FPCs) should spontaneously restart. We recommend that you wait until the warning no longer appears and then proceed with the switchover.

Options acquire—Attempt to become the master Routing Engine.

release—Request that the other Routing Engine become the master.

switch—Toggle mastership between Routing Engines.

The **acquire**, **release**, and **switch** options have the following suboptions:

no-confirm—(Optional) Do not request confirmation for the switch.

force—(Optional) Available only with the acquire option. Force the change to a new master Routing Engine.

Additional Information Because both Routing Engines are always running, the transition from one to the other as the master Routing Engine is immediate. However, the changeover interrupts communication to the System and Switch Board (SSB). The SSB takes several seconds to reinitialize the Flexible PIC Concentrators (FPCs) and restart the PICs. Interior gateway protocol (IGP) and BGP convergence times depend on the specific network environment.

By default, the Routing Engine in slot 0 (RE0) is the master and the Routing Engine in slot 1 (RE1) is the backup. To change the default master Routing Engine, include the **routing-engine** statement at the [**edit chassis redundancy**] hierarchy level in the configuration. For more information, see the *Junos OS System Basics Configuration Guide*

To have the backup Routing Engine become the master Routing Engine, use the **request chassis routing-engine master switch** command. If you use this command to change the master and then restart the chassis software for any reason, the master reverts to the default setting.

NOTE: Although the configurations on the two Routing Engines do not have to be the same and are not automatically synchronized, we recommend making both configurations the same.

Required Privilege Level maintenance

Related Documentation
- **show chassis routing-engine on page 653**

List of Sample Output
request chassis routing-engine master acquire on page 211
request chassis routing-engine master switch on page 211

Output Fields When you enter this command, you are provided feedback on the status of your request.

request chassis routing-engine master acquire

```
user@host> request chassis routing-engine master acquire

warning: Traffic will be interrupted while the PFE is re-initialized

warning: The other routing engine's file system could be corrupted

Reset other routing engine and become master ? [yes,no] (no)
```

request chassis routing-engine master switch

```
user@host> request chassis routing-engine master switch

warning: Traffic will be interrupted while the PFE is re-initialized
Toggle mastership between Routing Engines ? [yes,no] (no) yes

Resolving mastership...
Complete. The other Routing Engine becomes the master.
```

Switch mastership back to the local Routing Engine:

```
user@host> request chassis routing-engine master switch

warning: Traffic will be interrupted while the PFE is re-initialized
Toggle mastership between routing engines ? [yes,no] (no) yes

Resolving mastership...
Complete. The local routing engine becomes the master.
```

request system halt

Syntax

request system halt
<at *time*>
<both-routing-engines>
<other-routing-engine>
<in *minutes*>
<media (compact-flash | disk | removable-compact-flash | usb)>
<message "*text*">

Syntax (J-EX Series Switch)

request system halt
<all-members>
<at *time*>
<both-routing-engines>
<in *minutes*>
<local>
<media (external | internal)>
<member *member-id*>
<message "*text*">
<other-routing-engine>
<slice *slice*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Stop the router or switch software.

Options

none—Stop the router or switch software immediately.

all-members—(J-EX4200 switches only) (Optional) Halt all members of the Virtual Chassis configuration.

at *time* —(Optional) Time at which to stop the software, specified in one of the following ways:

- **now**—Stop the software immediately. This is the default.
- ***+minutes***—Number of minutes from now to stop the software.
- ***yymmddhhmm***—Absolute time at which to stop the software, specified as year, month, day, hour, and minute.
- ***hh:mm***—Absolute time on the current day at which to stop the software.

both-routing-engines—(Optional) Halt both Routing Engines at the same time.

local—(J-EX4200 switches only) (Optional) Halt the local Virtual Chassis member.

in *minutes*—(Optional) Number of minutes from now to stop the software. This option is an alias for the **at *+minutes*** option.

media (compact-flash | disk | removable-compact-flash | usb)—(Optional) Boot medium for next boot. (The options **removable-compact-flash** and **usb** pertain to J Series routers only.)

media (external | internal)—(J-EX Series switches only) (Optional) Halt the boot media:

- **external**—Halt the external mass storage device.
- **internal**—Halt the internal flash device.

member *member-id*—(J-EX4200 switches only) (Optional) Halt the specified member of the Virtual Chassis configuration. Replace ***member-id*** with a value from 0 through 9.

message "*text*"—(Optional) Message to display to all system users before stopping the software.

other-routing-engine—(Optional) Halt the other Routing Engine from which the command is issued. For example, if you issue the command from the master Routing Engine, the backup Routing Engine is halted. Similarly, if you issue the command from the backup Routing Engine, the master Routing Engine is halted.

slice *slice*—(J-EX Series switches only) (Optional) Halt a partition on the boot media. This option has the following suboptions:

- **1**—Halt partition 1.
- **2**—Halt partition 2.
- **alternate**—Reboot from the alternate partition.

Additional Information

NOTE: If you have a router or switch with two Routing Engines and you want to shut the power off to the router or switch or remove a Routing Engine, you must first halt the backup Routing Engine (if it has been upgraded), then halt the master Routing Engine. To halt a Routing Engine, issue the request system halt command. You can also halt both Routing Engines at the same time by issuing the request system halt both-routing-engines command.

Required Privilege Level

maintenance

List of Sample Output

request system halt on page 214
request system halt (in 2 Hours) on page 214
request system halt (Immediately) on page 214
request system halt (at 1:20 AM) on page 214

Output Fields

When you enter this command, you are provided feedback on the status of your request.

request system halt

```
user@host> request system halt
Halt the system ? [yes,no] (no) yes

*** FINAL System shutdown message from root@section2 ***
System going down IMMEDIATELY
Terminated
...
syncing disks... 11 8 done
The operating system has halted.
Please press any key to reboot.
```

request system halt (in 2 Hours)

The following example, which assumes that the time is 5 PM (1700), illustrates three different ways to request that the system stop 2 hours from now:

```
user@host> request system halt at +120
user@host> request system halt in 120
user@host> request system halt at 19:00
```

request system halt (Immediately)

```
user@host>  request system halt at now
```

request system halt (at 1:20 AM)

To stop the system at 1:20 AM, enter the following command. Because 1:20 AM is the next day, you must specify the absolute time.

```
user@host> request system halt at yymmdd120
request system halt at 120
Halt the system at 120? [yes,no] (no) yes
```

request system logout

Syntax request system logout (pid *pid* | terminal *terminal* | user *username*)
<all>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Log out users from the router or switch and the configuration database. If a user held the **configure exclusive** lock, this command clears the exclusive lock.

Options all—(Optional) Log out all sessions owned by a particular PID, terminal session, or user. (On a TX Matrix or TX Matrix Plus router, this command is broadcast to all chassis.)

pid *pid*—Log out the user session using the specified management process identifier (PID). The PID type must be management process.

terminal *terminal*—Log out the user for the specified terminal session.

user *username*—Log out the specified user.

Additional Information For information about using the **configure exclusive** command, see the *Junos OS System Basics Configuration Guide*.

Required Privilege Level configure

List of Sample Output **request system logout on page 215**

Output Fields When you enter this command, you are provided feedback on the status of your request.

request system logout

```
user@host> request system logout user tammy all
Connection closed by foreign host.
```

request system power-off

Syntax

request system power-off
<both-routing-engines>
<other-routing-engine>
<at *time*>
<in *minutes*>
<media (compact-flash | disk | removable-compact-flash | usb)>
<message "*text*">

Syntax (J-EX Series Switch)

request system power-off
<all-members>
<at *time*>
<both-routing-engines>
<in *minutes*>
<local>
<media (external | internal)>
<member *member-id*>
<message "*text*">
<other-routing-engine>
<slice *slice*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Power off the software.

Options

none—Power off the router or switch software immediately.

all-members—(J-EX4200 switches only) (Optional) Power off all members of the Virtual Chassis configuration.

at *time*—(Optional) Time at which to power off the software, specified in one of the following ways:

- **now**—Power off the software immediately. This is the default.
- ***+minutes***—Number of minutes from now to power off the software.
- ***yymmddhhmm***—Absolute time at which to power off the software, specified as year, month, day, hour, and minute.
- ***hh:mm***—Absolute time on the current day at which to power off the software.

both-routing-engines—(Optional) Power off both Routing Engines at the same time.

in *minutes*—(Optional) Number of minutes from now to power off the software. This option is an alias for the **at *+minutes*** option.

local—(J-EX4200 switches only) (Optional) Power off the local Virtual Chassis member.

media (compact-flash | disk | removable-compact-flash | usb)—(Optional) Boot medium for next boot. (The options **removable-compact-flash** and **usb** pertain to the J Series routers only.)

media (external | internal)—(J-EX Series switches only) (Optional) Power off the boot media:

- **external**—Power off the external mass storage device.
- **internal**—Power off the internal flash device.

member *member-id*—(J-EX4200 switches only) (Optional) Power off the specified member of the Virtual Chassis configuration. Replace ***member-id*** with a value from 0 through 9.

message "*text*"—(Optional) Message to display to all system users before powering off the software.

other-routing-engine—(Optional) Power off the other Routing Engine from which the command is issued. For example, if you issue the command from the master Routing Engine, the backup Routing Engine is halted. Similarly, if you issue the command from the backup Routing Engine, the master Routing Engine is halted.

slice *slice*—(J-EX Series switches only) (Optional) Power off a partition on the boot media. This option has the following suboptions:

- **1**—Power off partition 1.
- **2**—Power off partition 2.
- **alternate**—Reboot from the alternate partition.

Required Privilege Level maintenance

List of Sample Output **request system power-off on page 217**

Output Fields When you enter this command, you are provided feedback on the status of your request.

request system power-off

```
user@host> request system power-off message “This router will be powered off in 30 minutes.
Please save your data and log out immediately.”
warning: This command will not halt the other routing-engine.
If planning to switch off power, use the both-routing-engines option.
Power Off the system ? [yes,no] (no) yes

*** FINAL System shutdown message from remote@nutmeg ***
System going down IMMEDIATELY

This router will be powered off in 30 minutes. Please save your data and log out
immediately.

Shutdown NOW!
[pid 5177]
```

request system reboot

Syntax

request system reboot
<other-routing-engine>
<at *time*>
<in *minutes*>
<media (compact-flash | disk | removable-compact-flash | usb)>
<message "*text*">

Syntax (J-EX Series Switch)

request system reboot
<all-members>
<at *time*>
<in *minutes*>
<local>
<media (external | internal)>
<member *member-id*>
<message "*text*">
<other-routing-engine>
<slice *slice*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Reboot the software.

Options

none—Reboot the software immediately.

all-members—(J-EX4200 switches only) (Optional) Reboot all members of the Virtual Chassis configuration.

at *time*—(Optional) Time at which to reboot the software, specified in one of the following ways:

- **now**—Stop or reboot the software immediately. This is the default.
- ***+minutes***—Number of minutes from now to reboot the software.
- ***yymmddhhmm***—Absolute time at which to reboot the software, specified as year, month, day, hour, and minute.
- ***hh:mm***—Absolute time on the current day at which to stop the software, specified in 24-hour time.

in *minutes*—(Optional) Number of minutes from now to reboot the software. This option is an alias for the **at *+minutes*** option.

local—(J-EX4200 switches only) (Optional) Reboot the local Virtual Chassis member.

media (compact-flash | disk | removable-compact-flash | usb)—(Optional) Boot medium for next boot. (The options **removable-compact-flash** and **usb** pertain to the J Series routers only.)

media (external | internal)—(J-EX Series switches only) (Optional) Reboot the boot media:

- **external**—Reboot the external mass storage device.

- **internal**—Reboot the internal flash device.

member *member-id*—(J-EX4200 switches only) (Optional) Reboot the specified member of the Virtual Chassis configuration Replace ***member-id*** with a value from 0 through 9.

message "*text*"—(Optional) Message to display to all system users before stopping or rebooting the software.

other-routing-engine—(Optional) Reboot the other Routing Engine from which the command is issued. For example, if you issue the command from the master Routing Engine, the backup Routing Engine is rebooted. Similarly, if you issue the command from the backup Routing Engine, the master Routing Engine is rebooted.

slice *slice*—(J-EX Series switches only) (Optional) Reboot a partition on the boot media. This option has the following suboptions:

- **1**—Power off partition 1.
- **2**—Power off partition 2.
- **alternate**—Reboot from the alternate partition.

Additional Information

Reboot requests are recorded in the system log files, which you can view with the **show log** command (see **show log**). Also, the names of any running processes that are scheduled to be shut down are changed. You can view the process names with the **show system processes** command (see **show system processes**).

NOTE: To reboot a router that has two Routing Engines, reboot the backup Routing Engine (if you have upgraded it) first, and then reboot the master Routing Engine.

Required Privilege Level

maintenance

Related Documentation

- **clear system reboot on page 204**

List of Sample Output

request system reboot on page 219
request system reboot (at 2300) on page 220
request system reboot (in 2 Hours) on page 220
request system reboot (Immediately) on page 220
request system reboot (at 1:20 AM) on page 220

Output Fields

When you enter this command, you are provided feedback on the status of your request.

request system reboot

```
user@host> request system reboot
Reboot the system ? [yes,no] (no)
```

request system reboot (at 2300)

```
user@host> request system reboot at 2300 message ?Maintenance time!?
Reboot the system ? [yes,no] (no) yes

shutdown: [pid 186]
*** System shutdown message from root@berry.network.net ***
System going down at 23:00
```

request system reboot (in 2 Hours)

The following example, which assumes that the time is 5 PM (17:00), illustrates three different ways to request the system to reboot in two hours:

```
user@host> request system reboot at +120
user@host> request system reboot in 120
user@host> request system reboot at 19:00
```

request system reboot (Immediately)

```
user@host> request system reboot at now
```

request system reboot (at 1:20 AM)

To reboot the system at 1:20 AM, enter the following command. Because 1:20 AM is the next day, you must specify the absolute time.

```
user@host> request system reboot at 06060120
request system reboot at 120
Reboot the system at 120? [yes,no] (no) yes
```

request system reboot

Syntax

request system reboot
<all-members | local | member *member-id*>
<at *time*>
<in *minutes*>
<media (external | internal)>
<message "*text*">
<slice (1 | 2 | alternate)>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Reboot the Junos OS.

Reboot requests are recorded in the system log files, which you can view with the **show log** command. You can view the process names with the **show system processes** command.

Options

none—Reboots the software immediately.

all-members | **local** | **member** ***member-id***—(J-EX4200 switch only) (Optional) Specify which member of the Virtual Chassis to reboot:

- **all-members**—Reboots each switch that is a member of the Virtual Chassis.
- **local**—Reboots the local switch, meaning the switch you are logged into, only.
- **member** ***member-id***—Reboots the specified member switch of the Virtual Chassis.

at *time*—(Optional) Time at which to reboot the software, specified in one of the following ways:

- ***+minutes***—Number of minutes from now to reboot the software.
- ***hh:mm***—Absolute time on the current day at which to reboot the software, specified in 24-hour time.
- **now**—Stop or reboot the software immediately. This is the default.
- ***yymmddhhmm***—Absolute time at which to reboot the software, specified as year, month, day, hour, and minute.

in *minutes*—(Optional) Number of minutes from now to reboot the software. This option is an alias for the **at** ***+minutes*** option.

media (external | internal)—(Optional) Boot medium for the next boot. The external option reboots the switch using a software package stored on an external boot source, such as a USB flash drive. The internal option reboots the switch using a software package stored in an internal memory source.

message "*text*"—(Optional) Message to display to all system users before rebooting the software.

slice (1 | 2 | alternate)—(Optional) Reboot using the specified partition on the boot media. This option has the following suboptions:

- **1**—Reboot from partition 1.
- **2**—Reboot from partition 2.
- **alternate**—Reboot from the alternate partition, which is the partition that did not boot the switch at the last bootup.

Required Privilege Level maintenance

Related Documentation
- **clear system reboot on page 204**

Output Fields When you enter this command, you are provided feedback on the status of your request.

request system reboot

```
user@host> request system reboot
Reboot the system ? [yes,no] (no)
```

request system reboot (at 2300)

```
user@host> request system reboot at 2300 message ?Maintenance time!?
Reboot the system ? [yes,no] (no) yes

shutdown: [pid 186]
*** System shutdown message from root@berry.network.net ***
System going down at 23:00
```

request system reboot (in 2 Hours) The following example, which assumes that the time is 5 PM (17:00), illustrates three different ways to request the system to reboot in two hours:

```
user@host> request system reboot at +120
user@host> request system reboot in 120
user@host> request system reboot at 19:00
```

request system reboot (Immediately)

```
user@host> request system reboot at now
```

request system reboot (at 1:20 AM) To reboot the system at 1:20 AM, enter the following command. Because 1:20 AM is the next day, you must specify the absolute time.

```
user@host> request system reboot at 06060120
request system reboot at 120
Reboot the system at 120? [yes,no] (no) yes
```

request system scripts convert

Syntax **request system scripts convert (slax-to-xslt | xslt-to-slax) source *source/filename* destination *destination/<filename>***

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Convert an Extensible Stylesheet Language Transformations (XSLT) script to Stylesheet Language, Alternative syntaX (SLAX), or convert a SLAX script to XSLT.

Options destination *destination/<filename>*—Specify a destination for the converted file.

Optionally, you can specify a filename for the converted file. If you do not specify a filename, the software assigns one automatically. The default destination filename is the same as the source filename, except the file extension is altered. For example, the software converts a source file called **test.xsl** to **test.slax**. The software converts a source file called **test1.slax** to **test1.xsl**.

slax-to-xslt—Convert a SLAX script to XSLT.

source *source/filename*—Specify a source file that you want to convert.

xslt-to-slax—Convert an XSLT script to SLAX.

Required Privilege Level maintenance

List of Sample Output
request system scripts convert slax-to-xslt on page 223
request system scripts convert xslt-to-slax on page 223

Output Fields When you enter this command, you are provided feedback on the status of your request.

request system scripts convert slax-to-xslt

```
user@host> request system scripts convert slax-to-xslt source /var/db/scripts/op/script1.slax
destination /var/db/scripts/op
conversion complete
```

request system scripts convert xslt-to-slax

```
user@host> request system scripts convert xslt-to-slax source /var/db/scripts/commit/script1.xsl
destination /var/db/scripts/commit
conversion complete
```

request system scripts refresh-from commit

Syntax **request system scripts refresh-from commit file *file-name* url *url-path***

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Automatically download the initial Junos OS configuration and a set of standard commit scripts during a Junos XML management protocol/NETCONF session when a switch is brought up for the first time.

The Junos XML management protocol equivalent for this operational mode command is:

```
<request-script-refresh-from>
   <type>commit</type>
   <file>file-name></file>
   <URL>URL</URL>
</request-script-refresh-from>
```

Options file *file-name*—Name of the file to be downloaded.

url *url-path*—URL of the file to be downloaded.

Required Privilege Level maintenance

Related Documentation

- Understanding Automatic Refreshing of Scripts on J-EX Series Switches on page 323
- *Junos XML Management Protocol Guide* at http://www.juniper.net/techpubs/software/junos/
- *Junos OS NETCONF XML Management Protocol Guide* at http://www.juniper.net/techpubs/software/junos/

List of Sample Output **request system scripts refresh-from commit file config.txt url http://host1.juniper.net on page 224**

request system scripts refresh-from commit file config.txt url http://host1.juniper.net

```
user@switch> request system scripts refresh-from commit file config.txt url
http://host1.juniper.net
user@switch>
```

request system scripts refresh-from event

Syntax **request system scripts refresh-from event file *file-name* url *url-path***

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Automatically download the initial Junos OS configuration and a set of standard event scripts during a Junos XML management protocol/NETCONF session when a switch is brought up for the first time.

The Junos XML management protocol equivalent for this operational mode command is:

```
<request-script-refresh-from>
   <type>event</type>
   <file>file-name></file>
   <URL>URL</URL>
</request-script-refresh-from>
```

Options file *file-name*—Name of the file to be downloaded.

url *url-path*—URL of the file to be downloaded.

Required Privilege Level maintenance

Related Documentation

- Understanding Automatic Refreshing of Scripts on J-EX Series Switches on page 323
- *Junos OS Junos XML Management Protocol Guide* at http://www.juniper.net/techpubs/software/junos/
- *Junos OS NETCONF XML Management Protocol Guide* at http://www.juniper.net/techpubs/software/junos/

List of Sample Output **request system scripts refresh-from event file config.txt url http://host1.juniper.net on page 225**

request system scripts refresh-from event file config.txt url http://host1.juniper.net

```
user@switch> request system scripts refresh-from event file config.txt url http://host1.juniper.net
user@switch>
```

request system scripts refresh-from op

Syntax **request system scripts refresh-from op file *file-name* url *url-path***

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Automatically download the initial Junos OS configuration and a set of standard op scripts during a Junos XML management protocol/NETCONF session when a switch is brought up for the first time.

The Junos XML management protocol equivalent for this operational mode command is:

```
<request-script-refresh-from>
   <type>op</type>
   <file>file-name></file>
   <URL>URL</URL>
</request-script-refresh-from>
```

Options file *file-name*—Name of the file to be downloaded.

url *url-path*—URL of the file to be downloaded.

Required Privilege Level maintenance

Related Documentation

- Understanding Automatic Refreshing of Scripts on J-EX Series Switches on page 323
- *Junos OS Junos XML Management Protocol Guide* at http://www.juniper.net/techpubs/software/junos/
- *Junos OS NETCONF XML Management Protocol Guide* at http://www.juniper.net/techpubs/software/junos/

List of Sample Output **request system scripts refresh-from op file config.txt url http://host1.juniper.net on page 226**

request system scripts refresh-from op file config.txt url http://host1.juniper.net

```
user@switch> request system scripts refresh-from op file config.txt url http://host1.juniper.net
user@switch>
```

request system storage cleanup

Syntax request system storage cleanup <dry-run>

Syntax (J-EX Series Switch)
request system storage cleanup
<all-members>
<dry-run>
<local>
<member *member-id*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Free storage space on the router or switch by rotating log files and proposing a list of files for deletion. User input is required for file deletion.

Options
all-members—(J-EX4200 switches only) (Optional) Delete files on all members of the Virtual Chassis configuration.

dry-run—(Optional) List files proposed for deletion (without deleting them).

local—(J-EX4200 switches only) (Optional) Delete files on the local Virtual Chassis member.

member *member-id*—(J-EX4200 switches only) (Optional) Delete files on the specified member of the Virtual Chassis configuration. Replace ***member-id*** with a value from 0 through 9.

Additional Information If logging is configured and being used, the **dry-run** option will rotate the log files. In that case, the output displays the message "Currently rotating log files, please wait." If no logging is currently underway, the output displays only a list of files to delete.

Required Privilege Level maintenance

List of Sample Output
request system storage cleanup dry-run on page 227
request system storage cleanup on page 228

Output Fields When you enter this command, you are provided feedback on the status of your request.

request system storage cleanup dry-run

```
user@host> request system storage cleanup dry-run
Currently rotating log files, please wait.
This operation can take up to a minute.

List of files to delete:

         Size Date         Name
  11.4K Mar  8 15:00 /var/log/messages.1.gz
  7245B Feb  5 15:00 /var/log/messages.3.gz
  11.8K Feb 22 13:00 /var/log/messages.2.gz
  3926B Mar 16 13:57 /var/log/messages.0.gz
  3962B Feb 22 12:47 /var/log/sampled.1.gz
  4146B Mar  8 12:20 /var/log/sampled.0.gz
  4708B Dec 21 11:39 /var/log/sampled.2.gz
```

```
  7068B Jan 16 18:00 /var/log/messages.4.gz
  13.7K Dec 27 22:00 /var/log/messages.5.gz
   890B Feb 22 17:22 /var/tmp/sampled.pkts
  65.8M Oct 26 09:10 /var/sw/pkg/jinstall-7.4R1.7-export-signed.tgz
  63.1M Oct 26 09:13 /var/sw/pkg/jbundle-7.4R1.7.tgz
```

request system storage cleanup

```
user@host> request system storage cleanup
Currently rotating log files, please wait.
This operation can take up to a minute.

List of files to delete:

         Size Date         Name
  11.4K Mar  8 15:00 /var/log/messages.1.gz
  7245B Feb  5 15:00 /var/log/messages.3.gz
  11.8K Feb 22 13:00 /var/log/messages.2.gz
  3926B Mar 16 13:57 /var/log/messages.0.gz
  11.6K Mar  8 15:00 /var/log/messages.5.gz
  7254B Feb  5 15:00 /var/log/messages.6.gz
  12.9K Feb 22 13:00 /var/log/messages.8.gz
  3726B Mar 16 13:57 /var/log/messages.7.gz
  3962B Feb 22 12:47 /var/log/sampled.1.gz
  4146B Mar  8 12:20 /var/log/sampled.0.gz
  4708B Dec 21 11:39 /var/log/sampled.2.gz
  7068B Jan 16 18:00 /var/log/messages.4.gz
  13.7K Dec 27 22:00 /var/log/messages.5.gz
   890B Feb 22 17:22 /var/tmp/sampled.pkts
  65.8M Oct 26 09:10 /var/sw/pkg/jinstall-7.4R1.7-export-signed.tgz
  63.1M Oct 26 09:13 /var/sw/pkg/jbundle-7.4R1.7.tgz

Delete these files ? [yes,no] (yes)
```

restart

Syntax

restart
<adaptive-services | audit-process | chassis-control | class-of-service | dhcp-service | diameter-service | disk-monitoring | dynamic-flow-capture | ecc-error-logging | event-processing | firewall | interface-control | ipsec-key-management | kernel-replication | l2-learning | l2tp-service | lacp | mib-process | pgcp-service | pgm | pic-services-logging | ppp | pppoe | protected-system-domain-service | redundancy-interface-process | remote-operations | root-system-domain-service | routing <logical-system *logical-system-name*> | sampling | service-deployment | services pgcp gateway *gateway-name* |sbc-configuration-process | snmp | usb-control| web-management>
<gracefully | immediately | soft>

Syntax (J-EX Series Switch)

restart
<autoinstallation | chassis-control | class-of-service | database-replication | dhcp | dhcp-service | diameter-service | dot1x-protocol | ethernet-link-fault-management | ethernet-switching | event-processing | firewall | general-authentication-service | interface-control | kernel-replication | l2-learning | lacp | license-service | link-management | lldpd-service | mib-process | mountd-service | multicast-snooping | pgm | redundancy-interface-process | remote-operations | routing | secure-neighbor-discovery | service-deployment | sflow-service | snmp | vrrp | web-management>
<gracefully | immediately | soft>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Restart a Junos OS process.

CAUTION: Never restart a software process unless instructed to do so by a customer support engineer. A restart might cause the router or switch to drop calls and interrupt transmission, resulting in possible loss of data.

Options

none—Same as **gracefully**.

adaptive-services—(Optional) Restart the configuration management process that manages the configuration for stateful firewall, Network Address Translation (NAT), intrusion detection services (IDS), and IP Security (IPsec) services on the Adaptive Services PIC.

audit-process—(Optional) Restart the RADIUS accounting process.

autoinstallation—(J-EX Series switch only) (Optional) Restart the autoinstallation process.

chassis-control—(Optional) Restart the chassis management process.

class-of-service—(Optional) Restart the class-of-service (CoS) process, which controls the router's or switch's CoS configuration.

database-replication—(J-EX Series switch only) (Optional) Restart the database replication process.

dhcp—(J-EX Series switch only) (Optional) Restart the software process for a Dynamic Host Configuration Protocol (DHCP) server. A DHCP server allocates network IP addresses and delivers configuration settings to client hosts without user intervention.

dhcp-service—(J-EX Series switch only) (Optional) Restart the Dynamic Host Configuration Protocol process.

diameter-service—(Optional) Restart the diameter process.

disk-monitoring—(Optional) Restart disk monitoring, which checks the health of the hard disk drive on the Routing Engine.

dot1x-protocol—(J-EX Series switch only) (Optional) Restart the port-based network access control process.

dynamic-flow-capture—(Optional) Restart the dynamic flow capture (DFC) process, which controls DFC configurations on Monitoring Services III PICs.

ecc-error-logging—(Optional) Restart the error checking and correcting (ECC) process, which logs ECC parity errors in memory on the Routing Engine.

ethernet-link-fault-management—(J-EX Series switch only) (Optional) Restart the Ethernet OAM link fault management process.

ethernet-switching—(J-EX Series switch only) (Optional) Restart the Ethernet switching process.

event-processing—(Optional) Restart the event process (eventd).

firewall—(Optional) Restart the firewall management process, which manages firewall configuration.

general-authentication-service—(J-EX Series switch only) (Optional) Restart the general authentication process.

gracefully—(Optional) Restart the software process.

immediately—(Optional) Immediately restart the software process.

interface-control—(Optional) Restart the interface process, which controls the router's or switch's physical interface devices and logical interfaces.

ipsec-key-management—(Optional) Restart the IPsec key management process.

kernel-replication—(Optional) Restart the kernel replication process, which replicates the state of the backup Routing Engine when graceful Routing Engine switchover is configured.

l2-learning—(Optional) Restart the Layer 2 address flooding and learning process.

lacp—(Optional) Restart the Link Aggregation Control Protocol process.

license-service—(J-EX Series switch only) (Optional) Restart the feature license management process.

lldpd-service—(J-EX Series switch only) (Optional) Restart the Link Layer Discovery Protocol process.

mib-process—(Optional) Restart the Management Information Base (MIB) II process, which provides the router's MIB II agent.

mountd-service—(J-EX Series switch only) (Optional) Restart the service for NFS mounts requests.

multicast-snooping—(J-EX Series switch only) (Optional) Restart the multicast snooping process.

pgcp-service—(Optional) Restart the pgcpd service process running on the Routing Engine. This option does not restart pgcpd processes running on mobile station PICs. To restart pgcpd processes running on mobile station PICs, use the **services pgcp gateway** option.

pgm—(Optional) Restart the process that implements the Pragmatic General Multicast (PGM) protocol for assisting in the reliable delivery of multicast packets.

pic-services-logging—(Optional) Restart the logging process for some PICs. With this process, also known as fsad (the file system access daemon), PICs send special logging information to the Routing Engine for archiving on the hard disk.

ppp—(Optional) Restart the Point-to-Point Protocol (PPP) process.

pppoe—(Optional) Restart the Point-to-Point Protocol over Ethernet (PPPoE) process.

protected-system-domain-service—(Optional) Restart the Protected System Domain (PSD) process.

redundancy-interface-process—(Optional) Restart the ASP redundancy process.

remote-operations—(Optional) Restart the remote operations process, which provides the ping and traceroute MIBs.

root-system-domain-service—(Optional) Restart the Root System Domain (RSD) service.

routing—(J-EX Series switch only) (Optional) Restart the routing protocol process.

routing <logical-system *logical-system-name*>—(Optional) Restart the routing protocol process, which controls the routing protocols that run on the router or switch and maintains the routing tables. Optionally, restart the routing protocol process for the specified logical system only.

sampling—(Optional) Restart the sampling process, which performs packet sampling and cflowd export.

secure-neighbor-discovery—(J-EX Series switch only) (Optional) Restart the secure Neighbor Discovery Protocol process.

service-deployment—(Optional) Restart the service deployment service process.

services pgcp gateway *gateway-name*—(Optional) Restart the pgcpd process for a specific BGF running on an MS-PIC. This option does not restart the pgcpd process running on the Routing Engine. To restart the pgcpd process on the Routing Engine, use the **pgcp-service** option.

sflow-service—(J-EX Series switch only) (Optional) Restart the flow sampling (sFlow technology) process.

snmp—(Optional) Restart the SNMP process, which provides the router's or switch's SNMP master agent.

soft—(Optional) Reread and reactivate the configuration without completely restarting the software processes. For example, BGP peers stay up and the routing table stays constant. Omitting this option results in a graceful restart of the software process.

vrrp—(J-EX Series switch only) (Optional) Restart the Virtual Router Redundancy Protocol process.

web-management—(J-EX Series switch only) (Optional) Restart the Web management process.

Required Privilege Level reset

Related Documentation
- Overview of Junos OS CLI Operational Mode Commands

List of Sample Output **restart interfaces on page 232**

Output Fields When you enter this command, you are provided feedback on the status of your request.

restart interfaces

```
user@host> restart interfaces
interfaces process terminated
interfaces process restarted
```

set chassis display message

Syntax set chassis display message "*message*"
<permanent>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display or stop a text message on the craft interface display, which is on the front of the router, or on the LCD panel display on the switch. The craft interface alternates the display of text messages with standard craft interface messages, switching between messages every 2 seconds. By default, on both the router and the switch, the text message is displayed for 5 minutes. The craft interface display has four 20-character lines. The LCD panel display has two 16-character lines, and text messages appear only on the second line.

Options

"*message*"—Message to display. On the craft interface display, if the message is longer than 20 characters, it wraps onto the next line. If a word does not fit on one line, the entire word moves down to the next line. Any portion of the message that does not fit on the display is truncated. An empty pair of quotation marks (“ ”) deletes the text message from the craft interface display. On the LCD panel, display, the message is limited to 16 characters.

fpc-slot *slot-number*—(J-EX4200 switches only) On the router, display the text message on the craft interface for a specific Flexible PIC Concentrator (FPC). Replace ***slot-number*** with a value from **0** through **31**. On the switch, display the text message for a specific member of a virtual chassis, where **fpc-slot *slot-number*** corresponds to the member ID. Replace ***slot-number*** with a value from 0 through 9.

permanent—(Optional) Display a text message on the craft interface display or LCD panel display permanently.

Required Privilege Level clear

Related Documentation

- Configuring the LCD Panel on J-EX Series Switches (CLI Procedure) on page 166
- **clear chassis display message on page 202**
- show chassis craft-interface

List of Sample Output

set chassis display message (Creating) on page 233
set chassis display message (Deleting) on page 234

Output Fields See show chassis craft-interface for an explanation of output fields.

set chassis display message (Creating) The following example shows how to set the display message and verify the result:

```
user@host> set chassis display message "NOC contact Dusty (888) 555-1234"
message sent

user@host> show chassis craft-interface
```

```
Red alarm:     LED off, relay off
Yellow alarm:  LED off, relay off
Host OK LED:   On
Host fail LED: Off
FPCs     0  1  2  3  4  5  6  7
-------------------------------
Green  ..  *..  *  *.
Red    ........
LCD screen:
     +--------------------+
     |NOC contact Dusty   |
     |(888) 555-1234      |
     +--------------------+
```

set chassis display message (Deleting)

The following example shows how to delete the display message and verify that the message is removed:

```
user@host> set chassis display message ""
message sent


user@host> show chassis craft-interface
Red alarm:     LED off, relay off
Yellow alarm:  LED off, relay off
Host OK LED:   On
Host fail LED: Off
FPCs     0  1  2  3  4  5  6  7
-------------------------------
Green  ..  *..  *  *.
Red    ........
LCD screen:
     +--------------------+
     |host                |
     |Up: 0+17:05:47      |
     |                    |
     |Temperature OK      |
     +--------------------+
```

set date

Syntax set date (*date-time* | ntp <*ntp-server*> <source-address *source-address*>)

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set the date and time.

Options *date-time*—Date and time. Enter this string inside quotation marks.

ntp—Use a Network Time Protocol (NTP) server to synchronize the current date and time setting on the router or switch.

ntp-server—(Optional) Specify the IP address of one or more NTP servers.

source-address *source-address*—(Optional) Specify the source address that the router or switch uses to contact the remote NTP server.

Required Privilege Level view

Related Documentation

- **show cli on page 147**

List of Sample Output **set date on page 235**

Output Fields When you enter this command, you are provided feedback on the status of your request.

set date

```
user@host> set date ntp
21 Apr 17:22:02 ntpdate[3867]: step time server 172.17.27.46 offset 8.759252 sec
```

show chassis firmware

Syntax show chassis firmware

Release Information Command introduced in Junos OS Release 10.2 for J-EX Series switches.

Description On the routers, display the version levels of the firmware running on the System Control Board (SCB), Switching and Forwarding Module (SFM), System and Switch Board (SSB), Forwarding Engine Board (FEB), and Flexible PIC Concentrators (FPCs). On a TX Matrix Plus router, display the version levels of the firmware running on the FPCs and the Switch Processor Mezzanine Board (SPMBs).

On J-EX4200 switches, display the version levels of the firmware running on the switch. On a J-EX8208 switch, display the version levels of the firmware running on the Switch Fabric and Routing Engine (SRE) modules and on the line cards (shown as FPCs). On a J-EX8216 switch, display the version levels of the firmware running on the Routing Engine (RE) modules and on the line cards (shown as FPCs).

Options none—Display the version levels of the firmware running. For a J-EX4200 switch that is a member of a Virtual Chassis, display version levels for all members.

Required Privilege Level view

List of Sample Output **show chassis firmware (J-EX8200 Switch) on page 236**

Output Fields Table 33 on page 236 lists the output fields for the **show chassis firmware** command. Output fields are listed in the approximate order in which they appear.

Table 33: show chassis firmware Output Fields

Field Name	Field Description
Part	Chassis part name.
Type	Type of firmware: On routers: **ROM** or **O/S.** On switches: **uboot** or **loader**.
Version	Version of firmware running on the chassis part.

show chassis firmware (J-EX8200 Switch)

```
user@host> show chassis firmware

Part                    Type       Version
FPC 0                   U-Boot     U-Boot 1.1.6 (Mar 25 2009 - 06:13:12) 2.4.0
                        loader     FreeBSD/PowerPC U-Boot bootstrap loader 2.2
FPC 3                   U-Boot     U-Boot 1.1.6 (Dec  4 2009 - 13:17:34) 3.1.0
                        loader     FreeBSD/PowerPC U-Boot bootstrap loader 2.2
FPC 5                   U-Boot     U-Boot 1.1.6 (Mar 25 2009 - 06:13:12) 2.4.0
                        loader     FreeBSD/PowerPC U-Boot bootstrap loader 2.2
FPC 7                   U-Boot     U-Boot 1.1.6 (Feb  6 2009 - 05:31:46) 2.4.0
                        loader     FreeBSD/PowerPC U-Boot bootstrap loader 2.2
Routing Engine 0        U-Boot     U-Boot 1.1.6 (Mar 25 2009 - 06:13:12) 2.4.0
```

```
                                    FreeBSD/PowerPC U-Boot bootstrap loader 2.2
                    loader
Routing Engine 1    U-Boot          U-Boot 1.1.6 (Mar 25 2009 - 06:13:12) 2.4.0
                    loader          FreeBSD/PowerPC U-Boot bootstrap loader 2.2
```

show chassis lcd

Syntax

show chassis lcd
<fpc-slot <*fpc-slot-number*>>
<menu <(all-members | local | member *member-id*)>>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.
menu option introduced in Junos OS Release 10.2 for J-EX Series switches.

Description

Display the information that appears on the LCD panel of J-EX4200 and J-EX8200 switches. Display the status of the currently selected port parameter of the Status LED for each network port on the switch or line card.

Options

none—Display the information that appears on the LCD panel (for J-EX4200 switches configured as a Virtual Chassis, display the information for all Virtual Chassis members). Display the status of the currently selected port parameter of the Status LED for each network port.

fpc-slot <*fpc-slot-number*>—(Optional) Display the information as follows:

- For the standalone J-EX4200 switch (***fpc-slot-number*** equals 0)
- For all J-EX4200 switches in a Virtual Chassis (**fpc-slot** with no ***fpc-slot-number*** value specified)
- For a specific Virtual Chassis member (***fpc-slot-number*** equals member ID value)
- For the line card in the specified slot on a J-EX8200 switch (***fpc-slot-number*** equals slot number)

menu—(Optional) Display the names of the menus and menu options that are currently enabled on the LCD panel.

menu all-members—(J-EX4200 switches only) (Optional) Display the names of the menus and menu options that are currently enabled on the LCD panel for all Virtual Chassis members.

menu local—(J-EX4200 switches only) (Optional) Display the names of the menus and menu options that are currently enabled on the LCD panel for the Virtual Chassis member from which you issued the command.

menu member *member-id*—(J-EX4200 switches only) (Optional) Display the names of the menus and menu options that are currently enabled on the LCD panel for the specified Virtual Chassis member.

Required Privilege Level

view

Related Documentation

- LCD Panel in J-EX4200 Switches
- LCD Panel in a J-EX8200 Switch
- Configuring the LCD Panel on J-EX Series Switches (CLI Procedure) on page 166

List of Sample Output

Output Fields

Table 34 on page 239 lists the output fields for the **show chassis lcd** command. Output fields are listed in the approximate order in which they appear.

Table 34: show chassis lcd Output Fields

Field Name	Field Description
Front panel contents for slot **Front panel contents** (J-EX8200 switches)	FPC slot number of the switch whose content is being displayed. The number is always **0**, except for J-EX4200 switches in a Virtual Chassis, where it is the member ID value. On J-EX8200 switches, no slot number is displayed.
LCD screen	The first line displays the hostname (for Virtual Chassis members, displays the member ID, the current role, and hostname; for J-EX8200 switches, displays **RE** and the hostname). The second line displays the currently selected port parameter of the Status LED and the alarms counter. The Status LED port parameters are: • **ADM**—Administrative • **SPD**—Speed • **DPX**—Duplex • **POE**—Power over Ethernet (J-EX4200 switches only)
LEDs status	Current state of the Alarms, Status, and Master LEDs (chassis status LEDs).
Interface	Names of the interfaces on the switch.
LED (ADM/SPD/DPX/POE)	State of the currently selected port parameter of the Status LED for the interface. The Status LED port parameters are: NOTE: J-EX8200 switches do not have the POE port parameter. • **ADM**—Administrative • **SPD**—Speed • **DPX**—Duplex • **POE**—Power over Ethernet
fpc*x*	On J-EX4200 switches, member ID of the Virtual Chassis member whose LCD menu is displayed.

show chassis lcd (Two-Member Virtual Chassis)

```
user@switch> show chassis lcd
Front panel contents for slot: 0
--------------------------------
LCD screen:
```

```
    00:BK switch1
    LED:SPD ALARM 00
LEDs status:
    Alarms LED: Off
    Status LED: Green
    Master LED: Off
Interface       LED(ADM/SPD/DPX/POE)
-------------------------------------
ge-0/0/0        Off
ge-0/0/1        Off
ge-0/0/2        Off
ge-0/0/3        Off
ge-0/0/4        Off
ge-0/0/5        Off
ge-0/0/6        Off
ge-0/0/7        Off
ge-0/0/8        Off
ge-0/0/9        Off
ge-0/0/10       Off
ge-0/0/11       Off
ge-0/0/12       Off
ge-0/0/13       Off
ge-0/0/14       Off
ge-0/0/15       Off
ge-0/0/16       Off
ge-0/0/17       Off
ge-0/0/18       Off
ge-0/0/19       Off
ge-0/0/20       Off
ge-0/0/21       Off
ge-0/0/22       Off
ge-0/0/23       Off
Front panel contents for slot: 1
---------------------------------
LCD screen:
    01:RE switch2
    LED:SPD ALARM 01
LEDs status:
    Alarms LED: Yellow
    Status LED: Green
    Master LED: Green
Interface       LED(ADM/SPD/DPX/POE)
-------------------------------------
ge-1/0/0        Off
ge-1/0/1        Off
ge-1/0/2        Off
ge-1/0/3        Off
ge-1/0/4        Off
ge-1/0/5        Off
ge-1/0/6        Off
ge-1/0/7        Off
ge-1/0/8        Off
ge-1/0/9        Off
ge-1/0/10       Off
ge-1/0/11       Off
ge-1/0/12       Off
ge-1/0/13       Off
ge-1/0/14       Off
ge-1/0/15       Off
ge-1/0/16       Off
ge-1/0/17       Off
```

```
ge-1/0/18       Off
ge-1/0/19       Off
ge-1/0/20       Off
ge-1/0/21       Off
ge-1/0/22       Off
ge-1/0/23       Off
```

The output for the **show chassis lcd fpc-slot** command is the same as the output for the **show chassis lcd** command.

show chassis lcd fpc-slot 1 (Virtual Chassis)

```
user@switch> show chassis lcd fpc-slot 1
Front panel contents for slot: 1
---------------------------------
LCD screen:
    01:RE switch2
    LED:SPD ALARM 01
LEDs status:
    Alarms LED: Yellow
    Status LED: Green
    Master LED: Green
Interface       LED(ADM/SPD/DPX/POE)
------------------------------------
ge-1/0/0        Off
ge-1/0/1        Off
ge-1/0/2        Off
ge-1/0/3        Off
ge-1/0/4        Off
ge-1/0/5        Off
ge-1/0/6        Off
ge-1/0/7        Off
ge-1/0/8        Off
ge-1/0/9        Off
ge-1/0/10       Off
ge-1/0/11       Off
ge-1/0/12       Off
ge-1/0/13       Off
ge-1/0/14       Off
ge-1/0/15       Off
ge-1/0/16       Off
ge-1/0/17       Off
ge-1/0/18       Off
ge-1/0/19       Off
ge-1/0/20       Off
ge-1/0/21       Off
ge-1/0/22       Off
ge-1/0/23       Off
```

show chassis lcd (J-EX8200)

```
show chassis lcd
Front panel contents:
---------------------
LCD screen:
    RE st-8200-r
    LED:ADM ALARM 01
LEDs status:
    Alarms LED: Yellow
    Status LED: Yellow
    Master LED: Green
Interface       LED(ADM/SPD/DPX)
----------------------------------------
```

```
ge-0/0/0        Off
ge-0/0/1        Off
ge-0/0/2        Off
ge-0/0/3        Off
ge-0/0/4        Off
ge-0/0/5        Off
ge-0/0/6        Off
ge-0/0/7        Off
ge-0/0/8        Off
ge-0/0/9        Off
ge-0/0/10       Off
ge-0/0/11       Off
ge-0/0/12       Off
ge-0/0/13       Off
ge-0/0/14       Off
ge-0/0/15       Off
ge-0/0/16       Off
ge-0/0/17       Off
ge-0/0/18       Off
ge-0/0/19       Off
ge-0/0/20       Off
ge-0/0/21       Off
ge-0/0/22       Off
ge-0/0/23       Off
ge-0/0/24       Off
ge-0/0/25       Off
ge-0/0/26       Off
ge-0/0/27       Off
ge-0/0/28       Off
ge-0/0/29       Off
ge-0/0/30       Off
ge-0/0/31       Off
ge-0/0/32       Off
ge-0/0/33       Off
ge-0/0/34       Off
ge-0/0/35       Off
ge-0/0/36       Off
ge-0/0/37       Off
ge-0/0/38       Off
ge-0/0/39       Off
ge-0/0/40       Off
ge-0/0/41       Off
ge-0/0/42       Off
ge-0/0/43       Off
ge-0/0/44       Off
ge-0/0/45       Off
ge-0/0/46       Off
ge-0/0/47       Off
xe-2/0/0        Off
xe-2/0/1        Off
xe-2/0/2        Off
xe-2/0/3        Off
xe-2/0/4        Off
xe-2/0/5        Off
xe-2/0/6        Off
xe-2/0/7        Off
xe-3/0/0        Off
xe-3/0/1        Off
xe-3/0/2        Off
xe-3/0/3        Off
xe-3/0/4        Off
```

```
xe-3/0/5        Off
xe-3/0/6        Off
xe-3/0/7        Off
xe-5/0/0        Off
xe-5/0/1        Off
xe-5/0/2        Off
xe-5/0/3        Off
xe-5/0/4        Off
xe-5/0/5        Off
xe-5/0/6        On
xe-5/0/7        On
xe-7/0/5        Off
```

show chassis lcd fpc-slot 2 (J-EX8200)

```
show chassis lcd fpc-slot 2

Interface       LED(ADM/SPD/DPX)
----------------------------------------
xe-2/0/0        Off
xe-2/0/1        Off
xe-2/0/2        Off
xe-2/0/3        Off
xe-2/0/4        Off
xe-2/0/5        Off
xe-2/0/6        Off
xe-2/0/7        Off
```

show chassis lcd menu (J-EX4200)

```
user@switch> show chassis lcd menu
fpc0:
--------------------------------------------------------------------------------
status-menu
status-menu vcp-status
status-menu power-status
status-menu environ-menu
status-menu show-version
maintenance-menu
maintenance-menu halt-menu
maintenance-menu system-reboot
maintenance-menu rescue-config
maintenance-menu vc-uplink-config
maintenance-menu factory-default
```

On a J-EX4200 switch in a Virtual Chassis, the output for the **show chassis lcd menu all-members** command is the same as the output for the **show chassis lcd menu** command.

show chassis lcd menu (J-EX8200)

```
user@switch> show chassis lcd menu
status-menu
status-menu sf-status1-menu
status-menu sf-status2-menu
status-menu psu-status1-menu
status-menu psu-status2-menu
status-menu environ-menu
status-menu show-version
maintenance-menu
maintenance-menu halt-menu
maintenance-menu system-reboot
maintenance-menu rescue-config
maintenance-menu factory-default
```

show configuration

Syntax show configuration
<*statement-path*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display the configuration that currently is running on the router or switch, which is the last committed configuration.

Options none—Display the entire configuration.

statement-path—(Optional) Display one of the following hierarchies in a configuration. (Each ***statement-path*** option has additional suboptions not described here. See the appropriate configuration guide or J-EX Series switch documentation for more information.)

- **access**—Network access configuration.
- **access-profile**—Access profile configuration.
- **accounting-options**—Accounting data configuration.
- **applications**—Applications defined by protocol characteristics.
- **apply-groups**—Groups from which configuration data is inherited.
- **chassis**—Chassis configuration.
- **chassis network-services**—Current running mode.
- **class-of-service**—Class-of-service configuration.
- **diameter**—Diameter base protocol layer configuration.
- **ethernet-switching-options**—(J-EX Series switch only) Ethernet switching configuration.
- **event-options**—Event processing configuration.
- **firewall**—Firewall configuration.
- **forwarding-options**—Options that control packet sampling.
- **groups**—Configuration groups.
- **interfaces**—Interface configuration.
- **jsrc**—JSRC partition configuration.
- **jsrc-partition**—JSRC partition configuration.
- **logical-systems**—Logical system configuration.
- **poe**—(J-EX Series switch only) Power over Ethernet configuration.
- **policy-options**—Routing policy option configuration.
- **protocols**—Routing protocol configuration.

- **routing-instances**—Routing instance configuration.
- **routing-options**—Protocol-independent routing option configuration.
- **security**—Security configuration.
- **services**—Service PIC applications configuration.
- **snmp**—Simple Network Management Protocol configuration.
- **system**—System parameters configuration.
- **virtual-chassis**—(J-EX Series switch only) Virtual Chassis configuration.
- **vlans**—(J-EX Series switch only) VLAN configuration.

Additional Information The portions of the configuration that you can view depend on the user class that you belong to and the corresponding permissions. If you do not have permission to view a portion of the configuration, the text **ACCESS-DENIED** is substituted for that portion of the configuration. If you do not have permission to view authentication keys and passwords in the configuration, because the **secret** permission bit is not set for your user account, the text **SECRET-DATA** is substituted for that portion of the configuration. If an identifier in the configuration contains a space, the identifier is displayed in quotation marks.

Required Privilege Level view

Related Documentation
- Displaying the Current Junos OS Configuration
- Overview of Junos OS CLI Operational Mode Commands

List of Sample Output
show configuration on page 245
show configuration policy-options on page 246

Output Fields This command displays information about the current running configuration.

show configuration

```
user@host> show configuration
## Last commit: 2006-10-31 14:13:00 PST by alant version "8.2I0 [builder]"; ##
last changed: 2006-10-31 14:05:53 PST
system {
    host-name nestor;
    domain-name east.net;
    backup-router 192.1.1.254;
   time-zone America/Los_Angeles;
    default-address-selection;
    name-server {
        192.154.169.254;
        192.154.169.249;
        192.154.169.176;
    }
    services {
        telnet;
    }
    tacplus-server {
        1.2.3.4 {
            secret /* SECRET-DATA */;
```

```
            ...
        }
    }
}
interfaces {
    ...
}
protocols {
        isis {
            export "direct routes";
        }
}
policy-options {
    policy-statement "direct routes" {
        from protocol direct;
        then accept;
    }
}
```

show configuration policy-options

```
user@host> show configuration policy-options
policy-options {
    policy-statement "direct routes" {
        from protocol direct;
        then accept;
    }
}
```

show host

Syntax show host *hostname*

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display Domain Name System (DNS) hostname information.

Options *hostname*—Hostname or address.

Additional Information The **show host** command displays the raw data received from the DNS server.

Required Privilege Level view

List of Sample Output **show host on page 247**

show host

```
user@host> show host snark
snark.boojum.net has address 192.168.1.254

user@host> show host 192.168.1.254
Name: snark.boojum.net
Address: 192.168.1.254
Aliases:
```

show ntp associations

Syntax show ntp associations
<no-resolve>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display Network Time Protocol (NTP) peers and their state.

Options none—Display NTP peers and their state.

no-resolve—(Optional) Suppress symbolic addressing.

Required Privilege Level view

Related Documentation
- **show ntp status on page 250**

List of Sample Output **show ntp associations on page 249**

Output Fields Table 35 on page 248 describes the output fields for the **show ntp associations** command. Output fields are listed in the approximate order in which they appear.

Table 35: show ntp associations Output Fields

Field Name	Field Description
remote	Address or name of the remote NTP peer.
refid	Reference identifier of the remote peer. If the reference identifier is not known, this field shows a value of **0.0.0.0**.
st	Stratum of the remote peer.
t	Type of peer: **b** (broadcast), **l** (local), **m** (multicast), or **u** (unicast).
when	When the last packet from the peer was received.
poll	Polling interval, in seconds.
reach	Reachability register, in octal.
delay	Current estimated delay of the peer, in milliseconds.
offset	Current estimated offset of the peer, in milliseconds.
disp	Current estimated dispersion of the peer, in milliseconds.

Table 35: show ntp associations Output Fields *(continued)*

Field Name	Field Description
peer-name	Peer name and status of the peer in the clock selection process: • space—Discarded because of a high stratum value or failed sanity checks. • **x**—Designated "falseticker", by the intersection algorithm. • **.**—Culled from the end of the candidate list. • **-** —Discarded by the clustering algorithm. • **+**—Included in the final selection set. • **#**—Selected for synchronization, but the distance exceeds the maximum. • ***—Selected for synchronization. • **o**—Selected for synchronization, but the packets-per-second (pps) signal is in use.

show ntp associations

```
user@host> show ntp associations
     remote           refid      st t when poll reach   delay   offset    disp
==============================================================================
*wolfe-gw.junipe tick.ucla.edu    2 u   43   64  377     1.86    0.319    0.08
```

show ntp status

Syntax
show ntp status
<no-resolve>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display the values of internal variables returned by Network Time Protocol (NTP) peers.

Options
none—Display the values of internal variables returned by NTP peers.

no-resolve—(Optional) Suppress symbolic addressing.

Required Privilege Level
view

Related Documentation
- **show ntp associations on page 248**

List of Sample Output
show ntp status on page 250

show ntp status

```
user@host> show ntp status
status=0644 leap_none, sync_ntp, 4 events, event_peer/strat_chg,
version="ntpd 4.1.0-a Fri Jun 24 06:40:56 GMT 2005 (1)",
processor="i386", system="JUNOS7.4-20050624.0", leap=00, stratum=2,
precision=-28, rootdelay=6.849, rootdispersion=10.615, peer=38788,
refid=ntp-server.company-a.net,
reftime=c66705d9.06ee0f3c  Fri, Jun 24 2005 15:21:13.027, poll=6,
clock=c6670602.cf6db940  Fri, Jun 24 2005 15:21:54.810, state=4,
offset=0.205, frequency=75.911, jitter=0.396, stability=0.005
```

show system firmware

Syntax show system firmware
<compatibility>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description (J-EX8200 switches only) Display firmware information.

Options compatibility—(Optional) Display firmware compatibility information.

Required Privilege Level view

List of Sample Output
show system firmware on page 251
show system firmware compatibility on page 251

Output Fields Table 36 on page 251 lists the output fields for the show system firmware command. Output fields are listed in the approximate order in which they appear.

Table 36: show system firmware Output Fields

Field Name	Field Description
Part	Physical part on the router or switch affected by the firmware.
Type	Type of firmware on the router or switch.
Tag	Location of the firmware on the interface.
Current version	Firmware version on the affected router or switch parts.
Available version	New versions of firmware for upgrading or downgrading.
Status	Firmware condition on the router or switch.
Action	Whether you can upgrade or downgrade, or if no action is available (**none**).

show system firmware

```
user@host> show system firmware
Part             Type           Tag Current   Available Status
                                    version   version
FPC 0            ROM Monitor 0  0   6.4.10              OK
Routing Engine 0 RE BIOS        0   0                   OK
```

show system firmware compatibility

```
user@host> show system firmware compatibility
Part             Type           Tag Current   Available Action
                                    version   version
FPC 0            ROM Monitor 0  0   6.4.10              None
Routing Engine 0 RE BIOS        0   0                   None
```

show system reboot

Syntax
show system reboot
<both-routing-engines>

Syntax (J-EX Series Switch)
show system reboot
<all-members>
<both-routing-engines>
<local>
<member *member-id*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display pending system reboots or halts.

Options
none—Display pending reboots or halts on the active Routing Engine.

all-members—(J-EX4200 switches only) (Optional) Display halt or reboot request information for all members of the Virtual Chassis configuration.

both-routing-engines—(Systems with multiple Routing Engines) (Optional) Display halt or reboot request information on both Routing Engines.

local—(J-EX4200 switches only) (Optional) Display halt or reboot request information for the local Virtual Chassis member.

member *member-id*—(J-EX4200 switches only) (Optional) Display halt or reboot request information for the specified member of the Virtual Chassis configuration. Replace ***member-id*** with a value from 0 through 9.

Required Privilege Level
maintenance

List of Sample Output
show system reboot on page 252

show system reboot

```
user@host> show system reboot
reboot requested by root at Wed Feb 10 17:40:46 1999
[process id 17885]
```

show system snapshot

Syntax

show system snapshot
<all-members | local | member *member-id*>
<media (external | internal)>
<slice (1 | 2 | alternate)>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display the complete collection of files in a snapshot.

Options

none—Display the system snapshot on the alternate media, which is the media that does not have the software packages that last booted the switch.

all-members | local | member *member-id*—(J-EX4200 switch only) Display the snapshot in a Virtual Chassis configuration:

- **all-members**—Display the snapshot for each switch that is a member of the Virtual Chassis.
- **local**—Display the snapshot on the switch that you are currently logged into.
- **member *member-id***—Display the snapshot for the specified member switch of the Virtual Chassis.

media (external | internal)—(Optional) Display the destination media location for the snapshot. The **external** option specifies the snapshot on an external mass storage device, such as a USB flash drive. The **internal** option specifies the snapshot on an internal memory source, such as internal flash memory.

slice (1 | 2 | alternate)—Display the snapshot in a partition:

- **1**—Display the snapshot in partition 1.
- **2**—Display the snapshot in partition 2.
- **alternate**—Display the snapshot in the alternate partition, which is the partition that did not boot the switch at the last bootup.

Required Privilege Level

view

Related Documentation

- **request system snapshot on page 106**
- Creating a Snapshot and Using It to Boot a J-EX Series Switch on page 80
- Verifying That a System Snapshot Was Created on a J-EX Series Switch on page 88

show system snapshot media external

```
user@switch> show system snapshot media external
Information for snapshot on external (da1s1)
Creation date: Oct 13 20:23:23 2009
JUNOS version on snapshot:
jbase  : 10.0I20090726_0011_user
jcrypto-ex: 10.0I20090726_0011_user
```

```
jdocs-ex: 10.0I20090726_0011_user
jkernel-ex: 10.0I20090726_0011_user
jroute-ex: 10.0I20090726_0011_user
jswitch-ex: 10.0I20090726_0011_user
jweb-ex: 10.0I20090726_0011_user
jpfe-ex42x: 10.0I20090726_0011_user
```

show system software

Syntax
show system software
<detail>

Syntax (J-EX Series Switch)
show system software
<all-members>
<detail>
<local>
<member *member-id*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display the Junos OS extensions loaded on your router or switch.

Options
none—Display standard information about all loaded Junos OS extensions.

all-members—(J-EX4200 switches only) (Optional) Display the system software running on all members of the Virtual Chassis configuration.

detail—(Optional) Display detailed information about available Junos OS extensions.

local—(J-EX4200 switches only) (Optional) Display the system software running on the local Virtual Chassis member.

member *member-id*—(J-EX4200 switches only) (Optional) Display the system software running on the specified member of the Virtual Chassis configuration. Replace ***member-id*** with a value from 0 through 9.

scc—(Routing matrix only) (Optional) Display the system software running on a TX Matrix router (or switch-card chassis).

Required Privilege Level
maintenance

List of Sample Output
show system software on page 256

show system software

```
user@host> show system software
Information for jbase:

Comment:
JUNOS Base OS Software Suite [7.2R1.7]


Information for jcrypto:

Comment:
JUNOS Crypto Software Suite [7.2R1.7]
Information for jdocs:

Comment:
JUNOS Online Documentation [7.2R1.7]


Information for jkernel:
```

```
Comment:
JUNOS Kernel Software Suite [7.2R1.7]


Information for jpfe:

Comment:
JUNOS Packet Forwarding Engine Support (M20/M40) [7.2R1.7]


Information for jroute:

Comment:
JUNOS Routing Software Suite [7.2R1.7]


Information for junos:

Comment:
JUNOS Base OS boot [7.2R1.7]
```

show system storage

Syntax

show system storage
<detail>

Syntax (J-EX Series Switch)

show system storage
<detail>
<all-members>
<local>
<member *member-id*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display statistics about the amount of free disk space in the router's or switch's file systems.

Options

none—Display standard information about the amount of free disk space in the router's or switch's file systems.

detail—(Optional) Display detailed output.

all-members—(J-EX4200 switches only) (Optional) Display system storage statistics for all members of the Virtual Chassis configuration.

local—(J-EX4200 switches only) (Optional) Display system storage statistics for the local Virtual Chassis member.

member *member-id*—(J-EX4200 switches only) (Optional) Display system storage statistics for the specified member of the Virtual Chassis configuration. Replace ***member-id*** with a value from 0 through 9.

Required Privilege Level

view

List of Sample Output

show system storage on page 259

Output Fields

Table 37 on page 258 describes the output fields for the **show system storage** command. Output fields are listed in the approximate order in which they appear.

Table 37: show system storage Output Fields

Field Name	Field Description
Filesystem	Name of the file system.
Size	Size of the file system.
Used	Amount of space used in the file system.
Avail	Amount of space available in the file system.
Capacity	Percentage of the file system's space that is being used.

Table 37: show system storage Output Fields *(continued)*

Field Name	Field Description
Mounted on	Directory in which the file system is mounted.

show system storage

```
user@host> show system storage
Filesystem                Size       Used      Avail  Capacity   Mounted on
/dev/ad0s1a                77M        37M        34M       52%  /
devfs                      16K        16K         0B      100%  /dev/
/dev/vn0                   12M        12M         0B      100%  /packages/mnt/jbase
/dev/vn1                   39M        39M         0B      100%
/packages/mnt/jkernel-7.2R1.7
/dev/vn2                   12M        12M         0B      100%
/packages/mnt/jpfe-M40-7.2R1.7
/dev/vn3                  2.3M       2.3M         0B      100%
/packages/mnt/jdocs-7.2R1.7
/dev/vn4                   14M        14M         0B      100%
/packages/mnt/jroute-7.2R1.7
/dev/vn5                  4.5M       4.5M         0B      100%
/packages/mnt/jcrypto-7.2R1.7
mfs:172                   1.5G       4.0K       1.3G        0%  /tmp
/dev/ad0s1e                12M        20K        11M        0%  /config
procfs                    4.0K       4.0K         0B      100%  /proc
/dev/ad1s1f               9.4G       4.9G       3.7G       57%  /var
```

show system switchover

Syntax show system switchover

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display whether graceful Routing Engine switchover is configured, the state of the kernel replication (ready or synchronizing), any replication errors, and whether the primary and standby Routing Engines are using compatible versions of the kernel database.

NOTE: Issue the **show system switchover** command *only* on the backup Routing Engine. This command is *not* supported on the master Routing Engine. because the **kernel-replication** process daemon does not run on the master Routing Engine. This process runs only on the backup Routing Engine.

Options

Required Privilege Level view

List of Sample Output **show system switchover (Backup Routing Engine) on page 261**

Output Fields Table 38 on page 260 describes the output fields for the **show system switchover** command. Output fields are listed in the approximate order in which they appear.

Table 38: show system switchover Output Fields

Field Name	Field Description
Graceful switchover	Display graceful Routing Engine switchover status: • **On**—Indicates **graceful-switchover** is specified for the **routing-options** configuration command. • **Off**—Indicates **graceful-switchover** is not specified for the **routing-options** configuration command.
Configuration database	State of the configuration database: • **Ready**—Configuration database has synchronized. • **Synchronizing**—Configuration database is synchronizing. Displayed when there are updates within the last 5 seconds. • **Synchronize failed**—Configuration database synchronize process failed.
Kernel database	State of the kernel database: • **Ready**—Kernel database has synchronized. • **Synchronizing**—Kernel database is synchronizing. Displayed when there are updates within the last 5 seconds. • **Version incompatible**—The primary and standby Routing Engines are running incompatible kernel database versions. • **Replication error**—An error occurred when the state was replicated from the primary Routing Engine. Inspect **/var/log/ksyncd** for possible causes, or notify Dell Support (see "Requesting Technical Support" on page lxxi).

Table 38: show system switchover Output Fields *(continued)*

Field Name	Field Description
Peer state	Routing Engine peer state: • **Steady State**—Peer completed switchover transition. • **Peer Connected**—Peer in switchover transition.

show system switchover (Backup Routing Engine)

```
user@host> show system switchover
Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
```

show system uptime

Syntax show system uptime

Syntax (J-EX Series Switch)
show system uptime
<all-members>
<local>
<member *member-id*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display the current time and information about how long the router or switch, router or switch software, and routing protocols have been running.

Options
none—Show time since the system rebooted and processes started.

all-members—(J-EX4200 switches only) (Optional) Show time since the system rebooted and processes started on all members of the Virtual Chassis configuration.

local—(J-EX4200 switches only) (Optional) Show time since the system rebooted and processes started on the local Virtual Chassis member.

member *member-id*—(J-EX4200 switches only) (Optional) Show time since the system rebooted and processes started on the specified member of the Virtual Chassis configuration. Replace ***member-id*** with a value from 0 through 9.

Required Privilege Level view

List of Sample Output **show system uptime on page 263**

Output Fields Table 39 on page 262 describes the output fields for the **show system uptime** command. Output fields are listed in the approximate order in which they appear.

Table 39: show system uptime Output Fields

Field Name	Field Description
Current time	Current system time in UTC.
System booted	Date and time when the Routing Engine on the router or switch was last booted and how long it has been running.
Protocols started	Date and time when the routing protocols were last started and how long they have been running.
Last configured	Date and time when a configuration was last committed. Also shows name of user who issued the last **commit** command.
***time* and up**	Current time, in the local time zone, and how long the router or switch has been operational.
users	Number of users logged in to the router or router.

Table 39: show system uptime Output Fields *(continued)*

Field Name	Field Description
load averages	Load averages for the last 1 minute, 5 minutes, and 15 minutes.

show system uptime

```
user@host> show system uptime
Current time:      1998-10-13 19:45:47 UTC
System booted:     1998-10-12 20:51:41 UTC (22:54:06 ago)
Protocols started: 1998-10-13 19:33:45 UTC (00:12:02 ago)
Last configured:   1998-10-13 19:33:45 UTC (00:12:02 ago) by abc
12:45PM  up 22:54, 2 users, load averages: 0.07, 0.02, 0.01
```

show system users

Syntax show system users
<no-resolve>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description List information about the users who are currently logged in to the router or switch.

NOTE: The **show system users** command does not list information about the automated users that are currently logged in to the router or switch from a remote client application using Junos XML APIs, such as NETCONF. It only shows details of administrative users that are logged in to a router or switch using the CLI, J-Web, or an SSH client.

Options none—List information about the users who are currently logged in to the router or switch.

no-resolve—(Optional) Do not attempt to resolve IP addresses to hostnames.

Required Privilege Level view

List of Sample Output **show system users on page 265**

Output Fields Table 40 on page 264 describes the output fields for the **show system users** command. Output fields are listed in the approximate order in which they appear.

Table 40: show system users Output Fields

Field Name	Field Description
time **and up**	Current time, in the local time zone, and how long the router or switch has been operational.
users	Number of users logged in to the router or switch.
load averages	Load averages for the last 1 minute, 5 minutes, and 15 minutes.
USER	Username.
TTY	Terminal through which the user is logged in.
FROM	System from which the user has logged in. A hyphen indicates that the user is logged in through the console.
LOGIN@	Time when the user logged in.
IDLE	How long the user has been idle.
WHAT	Processes that the user is running.

show system users

```
user@host> show system users
 7:30PM  up 4 days,  2:26, 2 users, load averages: 0.07, 0.02, 0.01
USER     TTY FROM              LOGIN@  IDLE WHAT
root     d0  -                Fri05PM 4days -csh (csh)
blue   p0 level5.company.net 7:30PM     - cli
```

show system virtual-memory

Syntax show system virtual-memory

Syntax (J-EX Series Switch)
show system virtual-memory
<all-members>
<local>
<member *member-id*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display the usage of Junos OS kernel memory listed first by size of allocation and then by type of usage. Use **show system virtual-memory** for troubleshooting with Dell Support (see "Requesting Technical Support" on page lxxi).

Options none—Display kernel dynamic memory usage information.

all-members—(J-EX4200 switches only) (Optional) Display kernel dynamic memory usage information for all members of the Virtual Chassis configuration.

local—(J-EX4200 switches only) (Optional) Display kernel dynamic memory usage information for the local Virtual Chassis member.

member *member-id*—(J-EX4200 switches only) (Optional) Display kernel dynamic memory usage information for the specified member of the Virtual Chassis configuration. Replace ***member-id*** with a value from 0 through 9.

Additional Information

NOTE: The show system virtual-memory command with the | display XML pipe option displays XML output for the command in the parent tags: <vmstat-memstat-malloc>, <vmstat-memstat-zone>, <vmstat-sumstat>, <vmstat-intr>, and <vmstat-kernel-state> with each child element as a separate XML tag.

Required Privilege Level view

List of Sample Output
show system virtual-memory on page 268
show system virtual-memory | display xml on page 272

Output Fields Table 41 on page 267 lists the output fields for the **show system virtual-memory** command. Output fields are listed in the approximate order in which they appear.

Table 41: show system virtual-memory Output Fields

Field Name	Field Description
Memory statistics by bucket size	
Size	Memory block size (bytes). The kernel memory allocator appropriates blocks of memory whose size is exactly a power of 2.
In Use	Number of memory blocks of this size that are in use (bytes).
Free	Number of memory blocks of this size that are free (bytes).
Requests	Number of memory allocation requests made.
HighWater	Maximum value the free list can have. Once the system starts reclaiming physical memory, it continues until the free list is increased to this value.
Couldfree	Total number of times that the free elements for a bucket size exceed the high-water mark for that bucket size.
Memory usage type by bucket size	
Size	Memory block size (bytes).
Type(s)	Kernel modules that are using these memory blocks. For a definition of each type, refer to a FreeBSD book.
Memory statistics by type	
Type	Kernel module that is using dynamic memory.
InUse	Number of memory blocks used by this type. The number is rounded up.
MemUse	Amount of memory in use, in kilobytes (KB).
HighUse	Maximum memory ever used by this type.
Limit	Maximum memory that can be allocated to this type.
Requests	Total number of dynamic memory allocation requests this type has made.
Type Limit	Number of times requests were blocked for reaching the maximum limit.
Kern Limit	Number of times requests were blocked for kernel map.
Size(s)	Memory block sizes this type is using.
Memory Totals	
In Use	Total kernel dynamic memory in use (bytes, rounded up).
Free	Total kernel dynamic memory free (bytes, rounded up).

Table 41: show system virtual-memory Output Fields *(continued)*

Field Name	Field Description
Requests	Total number of memory allocation requests.
ITEM	Kernel module that is using memory.
Size	Memory block size (bytes).
LImit	Maximum memory that can be allocated to this type.
Used	Number of memory blocks used by this type. The number is rounded up.
Free	Number of memory blocks available to this type.
Requests	Total number of memory allocation requests this type has made.
interrupt	Timer events and scheduling interruptions.
total	Total number of interruptions for each type.
rate	Interruption rate.
Total	Total for all interruptions.

show system virtual-memory

```
user@host> show system virtual-memory
Memory statistics by bucket size
Size    In Use   Free   Requests  HighWater  Couldfree
  16       906    118     154876       1280          0
  32       455    313     209956        640          0
  64      4412    260      75380        320         20
 128      3200     32      19361        160         81
 256      1510     10       8844         80          4
 512       446      2       5085         40          0
  1K        18      2       5901         20          0
  2K      1128      2       4445         10       1368
  4K       185      1        456          5          0
  8K         5      1       2653          5          0
 16K       181      0        233          5          0
 32K         2      0       1848          5          0
 64K        20      0         22          5          0
128K         5      0          5          5          0
256K         2      0          2          5          0
512K         1      0          1          5          0

Memory usage type by bucket size
Size  Type(s)
  16  uc_devlist, nexusdev, iftable, temp, devbuf, atexit, COS, BPF,
        DEVFS mount, DEVFS node, vnodes, mount, pcb, soname, proc-args, kld,
        MD disk, rman, ATA generic, bus, sysctl, ippool, pfestat, ifstate,
        pfe_ipc, mkey, rtable, ifmaddr, ipfw, rnode
  32  atkbddev, dirrem, mkdir, diradd, freefile, freefrag, indirdep,
        bmsafemap, newblk, temp, devbuf, COS, vnodes, cluster_save buffer,
        pcb, soname, proc-args, sigio, kld, Gzip trees, taskqueue, SWAP,
```

```
         eventhandler, bus, sysctl, uidinfo, subproc, pgrp, pfestat, itable32,
         ifstate, pfe_ipc, mkey, rtable, ifmaddr, ipfw, rnode, rtnexthop
  64  isadev, iftable, MFS node, allocindir, allocdirect, pagedep, temp,
         devbuf, lockf, COS, NULLFS hash, DEVFS name, vnodes,
         cluster_save buffer, vfscache, pcb, soname, proc-args, file,
         AR driver, AD driver, Gzip trees, rman, eventhandler, bus, sysctl,
         subproc, pfestat, pic, ifstate, pfe_ipc, mkey, ifaddr, rtable, ipfw
 128  ZONE, freeblks, inodedep, temp, devbuf, zombie, COS, DEVFS node,
         vnodes, mount, vfscache, pcb, soname, proc-args, ttys, dev_t,
         timecounter, kld, Gzip trees, ISOFS node, bus, uidinfo, cred,
         session, pic, itable16, ifstate, pfe_ipc, rtable, ifstat, metrics,
         rtnexthop, iffamily
 256  iflogical, iftable, MFS node, FFS node, newblk, temp, devbuf,
         NFS daemon, vnodes, proc-args, kqueue, file desc, Gzip trees, bus,
         subproc, itable16, ifstate, pfe_ipc, sysctl, rtnexthop
 512  UFS mount, temp, devbuf, mount, BIO buffer, ptys, ttys, AR driver,
         Gzip trees, ISOFS mount, msg, ioctlops, ATA generic, bus, proc,
         pfestat, lr, ifstate, pfe_ipc, rtable, ipfw, ifstat, rtnexthop
  1K  iftable, temp, devbuf, NQNFS Lease, kqueue, kld, AD driver,
         Gzip trees, sem, MD disk, bus, ifstate, pfe_ipc, ipfw
  2K  uc_devlist, UFS mount, temp, devbuf, BIO buffer, pcb, AR driver,
         Gzip trees, ioctlops, bus, ipfw, ifstat, rcache
  4K  memdesc, iftable, UFS mount, temp, devbuf, kld, Gzip trees, sem, msg
  8K  temp, devbuf, syncache, Gzip trees
 16K  indirdep, temp, devbuf, shm, msg
 32K  pagedep, kld, Gzip trees
 64K  VM pgdata, devbuf, MSDOSFS mount
128K  UFS ihash, inodedep, NFS hash, kld, ISOFS mount
256K  mbuf, vfscache
512K  SWAP

Memory statistics by type                          Type  Kern
        Type  InUse MemUse HighUse  Limit Requests Limit Limit Size(s)
       isadev    13     1K      1K127753K       13    0     0  64
     atkbddev     2     1K      1K127753K        2    0     0  32
   uc_devlist    24     3K      3K127753K       24    0     0  16,2K
     nexusdev     3     1K      1K127753K        3    0     0  16
      memdesc     1     4K      4K127753K        1    0     0  4K
         mbuf     1   152K    152K127753K        1    0     0  256K
    iflogical     6     2K      2K127753K        6    0     0  256
      iftable    17     9K      9K127753K       18    0     0  16,64,256,1K,4K
         ZONE    15     2K      2K127753K       15    0     0  128
    VM pgdata     1    64K     64K127753K        1    0     0  64K
    UFS mount    12    26K     26K127753K       12    0     0  512,2K,4K
    UFS ihash     1   128K    128K127753K        1    0     0  128K
     MFS node     6     2K      3K127753K       35    0     0  64,256
     FFS node   906   227K    227K127753K     1352    0     0  256
       dirrem     0     0K      4K127753K      500    0     0  32
        mkdir     0     0K      1K127753K       38    0     0  32
       diradd     0     0K      6K127753K      521    0     0  32
     freefile     0     0K      4K127753K      374    0     0  32
     freeblks     0     0K      8K127753K      219    0     0  128
     freefrag     0     0K      1K127753K      193    0     0  32
   allocindir     0     0K     25K127753K     1518    0     0  64
     indirdep     0     0K     17K127753K       76    0     0  32,16K
  allocdirect     0     0K     10K127753K      760    0     0  64
    bmsafemap     0     0K      1K127753K       72    0     0  32
       newblk     1     1K      1K127753K     2279    0     0  32,256
     inodedep     1   128K    175K127753K     2367    0     0  128,128K
      pagedep     1    32K     33K127753K       47    0     0  64,32K
         temp  1239    92K     96K127753K     8364    0     0  16,32,64K
```

```
          devbuf  1413  5527K   5527K127753K     1535    0     0  16,32,64,128,256
           lockf    38     3K      3K127753K     2906    0     0  64
          atexit     1     1K      1K127753K        1    0     0  16
          zombie     0     0K      2K127753K     3850    0     0  128
        NFS hash     1   128K    128K127753K        1    0     0  128K
    NQNFS Lease      1     1K      1K127753K        1    0     0  1K
     NFS daemon      1     1K      1K127753K        1    0     0  256
       syncache      1     8K      8K127753K        1    0     0  8K
            COS    353    44K     44K127753K      353    0     0  16,32,64,128
            BPF    189     3K      3K127753K      189    0     0  16
  MSDOSFS mount      1    64K     64K127753K        1    0     0  64K
    NULLFS hash      1     1K      1K127753K        1    0     0  64
    DEVFS mount      2     1K      1K127753K        2    0     0  16
     DEVFS name    487    31K     31K127753K      487    0     0  64
     DEVFS node    471    58K     58K127753K      479    0     0  16,128
         vnodes     28     7K      7K127753K      429    0     0  16,32,64,128,256
          mount     15     8K      8K127753K       18    0     0  16,128,512
cluster_save buffer      0     0K      1K127753K       55    0     0  32,64
      vfscache  1898   376K    376K127753K     3228    0     0  64,128,256K
    BIO buffer    49    98K    398K127753K      495    0     0  512,2K
           pcb   159    16K     17K127753K      399    0     0  16,32,64,128,2K
        soname    82    10K     10K127753K    42847    0     0  16,32,64,128
     proc-args    57     2K      3K127753K     2105    0     0  16,32,64,128,256
          ptys    32    16K     16K127753K       32    0     0  512
          ttys   254    33K     33K127753K      522    0     0  128,512
        kqueue     5     3K      4K127753K       23    0     0  256,1K
         sigio     1     1K      1K127753K       27    0     0  32
          file   383    24K     24K127753K    16060    0     0  64
     file desc    76    19K     20K127753K     3968    0     0  256
           shm     1    12K     12K127753K        1    0     0  16K
         dev_t   286    36K     36K127753K      286    0     0  128
   timecounter    10     2K      2K127753K       10    0     0  128
           kld    11   117K    122K127753K       34    0     0  16,32,128,1K,4K
     AR driver     1     1K      3K127753K        5    0     0  64,512,2K
     AD driver     2     2K      3K127753K     2755    0     0  64,1K
     Gzip trees    0     0K     46K127753K   133848    0     0  32,64,128,256
     ISOFS node  1136   142K    142K127753K     1189    0     0  128
   ISOFS mount     9   132K    132K127753K       10    0     0  512,128K
           sem     3     6K      6K127753K        3    0     0  1K,4K
       MD disk     2     2K      2K127753K        2    0     0  16,1K
           msg     4    25K     25K127753K        4    0     0  512,4K,16K
          rman    59     4K      4K127753K      461    0     0  16,64
      ioctlops     0     0K      2K127753K      992    0     0  512,2K
     taskqueue     2     1K      1K127753K        2    0     0  32
          SWAP     2   413K    413K127753K        2    0     0  32,512K
   ATA generic     6     3K      3K127753K        6    0     0  16,512
  eventhandler    17     1K      1K127753K       17    0     0  32,64
           bus   340    30K     31K127753K      794    0     0  16,32,64,128,256
        sysctl     0     0K      1K127753K   130262    0     0  16,32,64
       uidinfo     4     1K      1K127753K       10    0     0  32,128
          cred    22     3K      3K127753K     3450    0     0  128
       subproc   156    10K     10K127753K     7882    0     0  32,64,256
          proc     2     1K      1K127753K        2    0     0  512
       session    12     2K      2K127753K       34    0     0  128
          pgrp    16     1K      1K127753K       45    0     0  32
        ippool     1     1K      1K127753K        1    0     0  16
       pfestat     0     0K      1K127753K    47349    0     0  16,32,64,512
           pic     5     1K      1K127753K        5    0     0  64,128
            lr     1     1K      1K127753K        1    0     0  512
     itable32   110     4K      4K127753K      110    0     0  32
     itable16   161    26K     26K127753K      161    0     0  128,256
```

```
      ifstate   694   159K    160K127753K     1735    0     0  16,32,64,128,1K
      pfe_ipc     0     0K      1K127753K    56218    0     0  16,32,64,128,1K
         mkey   250     4K      4K127753K      824    0     0  16,32,64
       ifaddr     9     1K      1K127753K        9    0     0  64
       sysctl     0     0K      1K127753K       30    0     0  256
       rtable    49     6K      6K127753K      307    0     0  16,32,64,128,512
      ifmaddr    22     1K      1K127753K       22    0     0  16,32
         ipfw    23    10K     10K127753K       48    0     0  16,32,64,512,2K
       ifstat   698   805K    805K127753K      698    0     0  128,512,2K
       rcache     4     8K      8K127753K        4    0     0  2K
        rnode    27     1K      1K127753K      285    0     0  16,32
      metrics     1     1K      1K127753K        3    0     0  128
    rtnexthop    57     9K      9K127753K      312    0     0  32,128,256,512
     iffamily    12     2K      2K127753K       12    0     0  128

Memory Totals:  In Use       Free    Requests
                 9311K        54K      489068

ITEM            SIZE     LIMIT    USED    FREE  REQUESTS
PIPE:            192,        0,      4,     81,      4422
SWAPMETA:        160,    95814,      0,      0,         0
unpcb:           160,        0,    114,     36,       279
ripcb:           192,    25330,      5,     37,         5
syncache:        128,    15359,      0,     64,         5
tcpcb:           576,    25330,     23,     12,        32
udpcb:           192,    25330,     14,     28,       255
socket:          256,    25330,    246,     26,       819
KNOTE:            96,        0,     27,     57,        71
NFSNODE:         352,        0,      0,      0,         0
NFSMOUNT:        544,        0,      0,      0,         0
VNODE:           224,        0,   2778,     43,      2778
NAMEI:          1024,        0,      0,      8,     40725
VMSPACE:         192,        0,     57,     71,      3906
PROC:            448,        0,     73,     17,      3923
DP fakepg:        64,        0,      0,      0,         0
PV ENTRY:         28,   499566,  44530, 152053,   1525141
MAP ENTRY:        48,        0,   1439,    134,    351075
KMAP ENTRY:       48,    35645,    179,    119,     10904
MAP:             108,        0,      7,      3,         7
VM OBJECT:        92,        0,   2575,    109,     66912

   792644 cpu context switches
  9863474 device interrupts
   286510 software interrupts
   390851 traps
  3596829 system calls
       16 kernel threads created
     3880  fork() calls
       27 vfork() calls
        0 rfork() calls
        0 swap pager pageins
        0 swap pager pages paged in
        0 swap pager pageouts
        0 swap pager pages paged out
      380 vnode pager pageins
      395 vnode pager pages paged in
      122 vnode pager pageouts
     1476 vnode pager pages paged out
        0 page daemon wakeups
        0 pages examined by the page daemon
      101 pages reactivated
```

```
   161722 copy-on-write faults
        0 copy-on-write optimized faults
    84623 zero fill pages zeroed
    83063 zero fill pages prezeroed
        7 intransit blocking page faults
   535606 total VM faults taken
        0 pages affected by kernel thread creation
   238254 pages affected by  fork()
     2535 pages affected by vfork()
        0 pages affected by rfork()
   283379 pages freed
        0 pages freed by daemon
   190091 pages freed by exiting processes
    17458 pages active
    29166 pages inactive
        0 pages in VM cache
    10395 pages wired down
   134610 pages free
     4096 bytes per page
   183419 total name lookups
          cache hits (90% pos + 7% neg) system 0% per-directory
          deletions 0%, falsehits 0%, toolong 0%

interrupt                   total       rate
ata0 irq14                 113338          3
mux irq7                   727643         21
fxp1 irq10                1178671         34
sio0 irq4                     833          0
clk irq0                 3439769         99
rtc irq8                 4403221        127
Total                    9863475        286
```

show system virtual-memory | display xml

```
user@host> show system virtual-memory | display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/10.2R1/junos">
    <system-virtual-memory-information>
        <vmstat-memstat-malloc>
            <memstat-name>CAM dev queue</memstat-name>
            <inuse>1</inuse>
            <memuse>1</memuse>
            <high-use>-</high-use>
            <memstat-req>1</memstat-req>
            <memstat-size>64</memstat-size>
            <memstat-name>entropy</memstat-name>
            <inuse>1024</inuse>
            <memuse>64</memuse>
            <high-use>-</high-use>
            <memstat-req>1024</memstat-req>
            <memstat-size>64</memstat-size>
            <memstat-name>linker</memstat-name>
            <inuse>481</inuse>
            <memuse>1871</memuse>
            <high-use>-</high-use>
            <memstat-req>1145</memstat-req>
            <memstat-size>16,32,64,4096,32768,131072</memstat-size>
            <memstat-name>lockf</memstat-name>
            <inuse>56</inuse>
            <memuse>4</memuse>
            <high-use>-</high-use>
            <memstat-req>5998</memstat-req>
            <memstat-size>64</memstat-size>
            <memstat-name>devbuf</memstat-name>
```

```
                <inuse>2094</inuse>
                <memuse>3877</memuse>
                <high-use>-</high-use>
                <memstat-req>2099</memstat-req>

<memstat-size>16,32,64,128,512,1024,4096,8192,16384,32768,65536,131072</memstat-size>

                <memstat-name>temp</memstat-name>
                <inuse>21</inuse>
                <memuse>66</memuse>
                <high-use>-</high-use>
                <memstat-req>3127</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,4096,8192,16384,32768,65536,131072</memstat-size>

                <memstat-name>ip6ndp</memstat-name>
                <inuse>0</inuse>
                <memuse>0</memuse>
                <high-use>-</high-use>
                <memstat-req>4</memstat-req>
                <memstat-size>64</memstat-size>
                <memstat-name>in6ifmulti</memstat-name>
                <inuse>1</inuse>
                <memuse>1</memuse>
                <high-use>-</high-use>
                <memstat-req>1</memstat-req>
                <memstat-size>64</memstat-size>
                <memstat-name>in6grentry</memstat-name>
                <inuse>1</inuse>
                <memuse>1</memuse>
                <high-use>-</high-use>
                <memstat-req>1</memstat-req>
                <memstat-size>64</memstat-size>
                <memstat-name>iflogical</memstat-name>
                <inuse>13</inuse>
                <memuse>3</memuse>
                <high-use>-</high-use>
                <memstat-req>13</memstat-req>
                <memstat-size>64,2048</memstat-size>
                <memstat-name>iffamily</memstat-name>
                <inuse>28</inuse>
                <memuse>4</memuse>
                <high-use>-</high-use>
                <memstat-req>28</memstat-req>
                <memstat-size>32,1024,2048</memstat-size>
                <memstat-name>rtnexthop</memstat-name>
                <inuse>127</inuse>
                <memuse>18</memuse>
                <high-use>-</high-use>
                <memstat-req>129</memstat-req>
                <memstat-size>32,256,512,1024,2048,4096</memstat-size>
                <memstat-name>metrics</memstat-name>
                <inuse>3</inuse>
                <memuse>1</memuse>
                <high-use>-</high-use>
                <memstat-req>5</memstat-req>
                <memstat-size>256</memstat-size>
                <memstat-name>inifmulti</memstat-name>
                <inuse>3</inuse>
                <memuse>1</memuse>
                <high-use>-</high-use>
```

```
                    <memstat-req>3</memstat-req>
                    <memstat-size>64</memstat-size>
                    <memstat-name>ingrentry</memstat-name>
                    <inuse>6</inuse>
                    <memuse>1</memuse>
                    <high-use>-</high-use>
                    <memstat-req>6</memstat-req>
                    <memstat-size>64</memstat-size>
                    <memstat-name>rnode</memstat-name>
                    <inuse>68</inuse>
                    <memuse>2</memuse>
                    <high-use>-</high-use>
                    <memstat-req>76</memstat-req>
                    <memstat-size>16,32</memstat-size>
                    <memstat-name>rcache</memstat-name>
                    <inuse>4</inuse>
                    <memuse>8</memuse>
                    <high-use>-</high-use>
                    <memstat-req>4</memstat-req>
                    <memstat-size>65536</memstat-size>
                    <memstat-name>ifdevice</memstat-name>
                    <inuse>4</inuse>
                    <memuse>1</memuse>
                    <high-use>-</high-use>
                    <memstat-req>4</memstat-req>
                    <memstat-size>16</memstat-size>
                    <memstat-name>ifstat</memstat-name>
                    <inuse>40</inuse>
                    <memuse>22</memuse>
                    <high-use>-</high-use>
                    <memstat-req>40</memstat-req>
                    <memstat-size>512,16384,32768</memstat-size>
                    <memstat-name>ipfw</memstat-name>
                    <inuse>42</inuse>
                    <memuse>23</memuse>
                    <high-use>-</high-use>
                    <memstat-req>91</memstat-req>

<memstat-size>16,32,64,128,256,512,1024,16384,32768,65536,131072</memstat-size>
                    <memstat-name>ifmaddr</memstat-name>
                    <inuse>103</inuse>
                    <memuse>3</memuse>
                    <high-use>-</high-use>
                    <memstat-req>103</memstat-req>
                    <memstat-size>16,32</memstat-size>
                    <memstat-name>rtable</memstat-name>
                    <inuse>129</inuse>
                    <memuse>14</memuse>
                    <high-use>-</high-use>
                    <memstat-req>139</memstat-req>
                    <memstat-size>16,32,64,128,1024,16384</memstat-size>
                    <memstat-name>sysctl</memstat-name>
                    <inuse>0</inuse>
                    <memuse>0</memuse>
                    <high-use>-</high-use>
                    <memstat-req>14847</memstat-req>
                    <memstat-size>16,32,64,4096,16384,32768</memstat-size>
                    <memstat-name>ifaddr</memstat-name>
                    <inuse>29</inuse>
                    <memuse>3</memuse>
                    <high-use>-</high-use>
```

```
                <memstat-req>29</memstat-req>
                <memstat-size>64,128</memstat-size>
                <memstat-name>mkey</memstat-name>
                <inuse>345</inuse>
                <memuse>6</memuse>
                <high-use>-</high-use>
                <memstat-req>2527</memstat-req>
                <memstat-size>16,128</memstat-size>
                <memstat-name>pfe_ipc</memstat-name>
                <inuse>0</inuse>
                <memuse>0</memuse>
                <high-use>-</high-use>
                <memstat-req>1422</memstat-req>

<memstat-size>16,32,64,128,512,1024,2048,8192,16384,32768,65536,131072</memstat-size>

                <memstat-name>ifstate</memstat-name>
                <inuse>594</inuse>
                <memuse>51</memuse>
                <high-use>-</high-use>
                <memstat-req>655</memstat-req>

<memstat-size>16,32,64,128,256,1024,2048,4096,16384,32768</memstat-size>
                <memstat-name>itable16</memstat-name>
                <inuse>276</inuse>
                <memuse>52</memuse>
                <high-use>-</high-use>
                <memstat-req>294</memstat-req>
                <memstat-size>1024,4096</memstat-size>
                <memstat-name>itable32</memstat-name>
                <inuse>160</inuse>
                <memuse>10</memuse>
                <high-use>-</high-use>
                <memstat-req>160</memstat-req>
                <memstat-size>64</memstat-size>
                <memstat-name>itable64</memstat-name>
                <inuse>2</inuse>
                <memuse>1</memuse>
                <high-use>-</high-use>
                <memstat-req>2</memstat-req>
                <memstat-size>128</memstat-size>
                <memstat-name>lr</memstat-name>
                <inuse>1</inuse>
                <memuse>1</memuse>
                <high-use>-</high-use>
                <memstat-req>1</memstat-req>
                <memstat-size>16384</memstat-size>
                <memstat-name>pic</memstat-name>
                <inuse>5</inuse>
                <memuse>1</memuse>
                <high-use>-</high-use>
                <memstat-req>5</memstat-req>
                <memstat-size>64,512</memstat-size>
                <memstat-name>pfestat</memstat-name>
                <inuse>0</inuse>
                <memuse>0</memuse>
                <high-use>-</high-use>
                <memstat-req>162</memstat-req>
                <memstat-size>16,32,128,256,16384</memstat-size>
                <memstat-name>gencfg</memstat-name>
                <inuse>224</inuse>
```

```
<memuse>56</memuse>
<high-use>-</high-use>
<memstat-req>540</memstat-req>
<memstat-size>16,32,64,256,512,32768,65536</memstat-size>
<memstat-name>jsr</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>-</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>idl</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>-</high-use>
<memstat-req>13</memstat-req>
<memstat-size>16,32,64,128,256,4096,16384,32768,131072</memstat-size>

<memstat-name>rtsmsg</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>-</high-use>
<memstat-req>2</memstat-req>
<memstat-size>131072</memstat-size>
<memstat-name>module</memstat-name>
<inuse>249</inuse>
<memuse>16</memuse>
<high-use>-</high-use>
<memstat-req>249</memstat-req>
<memstat-size>64,128</memstat-size>
<memstat-name>mtx_pool</memstat-name>
<inuse>1</inuse>
<memuse>8</memuse>
<high-use>-</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64,128</memstat-size>
<memstat-name>DEVFS3</memstat-name>
<inuse>109</inuse>
<memuse>12</memuse>
<high-use>-</high-use>
<memstat-req>117</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>DEVFS1</memstat-name>
<inuse>102</inuse>
<memuse>23</memuse>
<high-use>-</high-use>
<memstat-req>109</memstat-req>
<memstat-size>2048</memstat-size>
<memstat-name>pgrp</memstat-name>
<inuse>12</inuse>
<memuse>1</memuse>
<high-use>-</high-use>
<memstat-req>21</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>session</memstat-name>
<inuse>8</inuse>
<memuse>1</memuse>
<high-use>-</high-use>
<memstat-req>15</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>proc</memstat-name>
<inuse>2</inuse>
```

```
<memuse>1</memuse>
<high-use>-</high-use>
<memstat-req>2</memstat-req>
<memstat-size>16384</memstat-size>
<memstat-name>subproc</memstat-name>
<inuse>244</inuse>
<memuse>496</memuse>
<high-use>-</high-use>
<memstat-req>1522</memstat-req>
<memstat-size>2048,131072</memstat-size>
<memstat-name>cred</memstat-name>
<inuse>30</inuse>
<memuse>4</memuse>
<high-use>-</high-use>
<memstat-req>11409</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>plimit</memstat-name>
<inuse>17</inuse>
<memuse>4</memuse>
<high-use>-</high-use>
<memstat-req>133</memstat-req>
<memstat-size>2048</memstat-size>
<memstat-name>uidinfo</memstat-name>
<inuse>3</inuse>
<memuse>1</memuse>
<high-use>-</high-use>
<memstat-req>6</memstat-req>
<memstat-size>32,512</memstat-size>
<memstat-name>sysctloid</memstat-name>
<inuse>1117</inuse>
<memuse>34</memuse>
<high-use>-</high-use>
<memstat-req>1117</memstat-req>
<memstat-size>16,32,64</memstat-size>
<memstat-name>sysctltmp</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>-</high-use>
<memstat-req>743</memstat-req>
<memstat-size>16,32,64,1024</memstat-size>
<memstat-name>umtx</memstat-name>
<inuse>144</inuse>
<memuse>9</memuse>
<high-use>-</high-use>
<memstat-req>144</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>SWAP</memstat-name>
<inuse>2</inuse>
<memuse>209</memuse>
<high-use>-</high-use>
<memstat-req>2</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>bus</memstat-name>
<inuse>496</inuse>
<memuse>55</memuse>
<high-use>-</high-use>
<memstat-req>1196</memstat-req>
<memstat-size>16,32,64,128,32768</memstat-size>
<memstat-name>bus-sc</memstat-name>
<inuse>23</inuse>
<memuse>33</memuse>
```

```
                <high-use>-</high-use>
                <memstat-req>335</memstat-req>

<memstat-size>16,32,64,512,1024,2048,8192,16384,65536,131072</memstat-size>
                <memstat-name>devstat</memstat-name>
                <inuse>10</inuse>
                <memuse>21</memuse>
                <high-use>-</high-use>
                <memstat-req>10</memstat-req>
                <memstat-size>16,131072</memstat-size>
                <memstat-name>eventhandler</memstat-name>
                <inuse>35</inuse>
                <memuse>2</memuse>
                <high-use>-</high-use>
                <memstat-req>36</memstat-req>
                <memstat-size>32,128</memstat-size>
                <memstat-name>kobj</memstat-name>
                <inuse>93</inuse>
                <memuse>186</memuse>
                <high-use>-</high-use>
                <memstat-req>111</memstat-req>
                <memstat-size>65536</memstat-size>
                <memstat-name>DEVFS</memstat-name>
                <inuse>8</inuse>
                <memuse>1</memuse>
                <high-use>-</high-use>
                <memstat-req>9</memstat-req>
                <memstat-size>16,64</memstat-size>
                <memstat-name>rman</memstat-name>
                <inuse>71</inuse>
                <memuse>5</memuse>
                <high-use>-</high-use>
                <memstat-req>433</memstat-req>
                <memstat-size>16,32,64</memstat-size>
                <memstat-name>sbuf</memstat-name>
                <inuse>0</inuse>
                <memuse>0</memuse>
                <high-use>-</high-use>
                <memstat-req>522</memstat-req>
                <memstat-size>16,32,32768,131072</memstat-size>
                <memstat-name>NULLFS hash</memstat-name>
                <inuse>1</inuse>
                <memuse>1</memuse>
                <high-use>-</high-use>
                <memstat-req>1</memstat-req>
                <memstat-size>64</memstat-size>
                <memstat-name>taskqueue</memstat-name>
                <inuse>5</inuse>
                <memuse>1</memuse>
                <high-use>-</high-use>
                <memstat-req>5</memstat-req>
                <memstat-size>64</memstat-size>
                <memstat-name>turnstiles</memstat-name>
                <inuse>145</inuse>
                <memuse>10</memuse>
                <high-use>-</high-use>
                <memstat-req>145</memstat-req>
                <memstat-size>64</memstat-size>
                <memstat-name>Unitno</memstat-name>
                <inuse>8</inuse>
                <memuse>1</memuse>
```

```
<high-use>-</high-use>
<memstat-req>44</memstat-req>
<memstat-size>16,64</memstat-size>
<memstat-name>ioctlops</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>-</high-use>
<memstat-req>27622</memstat-req>
<memstat-size>16,64,8192,16384,131072</memstat-size>
<memstat-name>iov</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>-</high-use>
<memstat-req>18578</memstat-req>
<memstat-size>16,64,128,256,512,1024,2048,131072</memstat-size>
<memstat-name>msg</memstat-name>
<inuse>4</inuse>
<memuse>25</memuse>
<high-use>-</high-use>
<memstat-req>4</memstat-req>
<memstat-size>32768,131072</memstat-size>
<memstat-name>sem</memstat-name>
<inuse>4</inuse>
<memuse>7</memuse>
<high-use>-</high-use>
<memstat-req>4</memstat-req>
<memstat-size>16384,32768,131072</memstat-size>
<memstat-name>shm</memstat-name>
<inuse>9</inuse>
<memuse>20</memuse>
<high-use>-</high-use>
<memstat-req>14</memstat-req>
<memstat-size>32768</memstat-size>
<memstat-name>ttys</memstat-name>
<inuse>321</inuse>
<memuse>61</memuse>
<high-use>-</high-use>
<memstat-req>528</memstat-req>
<memstat-size>512,32768</memstat-size>
<memstat-name>ptys</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>-</high-use>
<memstat-req>1</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>mbuf_tag</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>-</high-use>
<memstat-req>23383</memstat-req>
<memstat-size>16</memstat-size>
<memstat-name>soname</memstat-name>
<inuse>115</inuse>
<memuse>12</memuse>
<high-use>-</high-use>
<memstat-req>24712</memstat-req>
<memstat-size>16,32,64,256</memstat-size>
<memstat-name>pcb</memstat-name>
<inuse>216</inuse>
<memuse>33</memuse>
<high-use>-</high-use>
```

```
                <memstat-req>484</memstat-req>

<memstat-size>16,32,64,128,1024,2048,4096,16384,32768,65536</memstat-size>
                <memstat-name>BIO buffer</memstat-name>
                <inuse>43</inuse>
                <memuse>86</memuse>
                <high-use>-</high-use>
                <memstat-req>405</memstat-req>
                <memstat-size>65536</memstat-size>
                <memstat-name>vfscache</memstat-name>
                <inuse>1</inuse>
                <memuse>256</memuse>
                <high-use>-</high-use>
                <memstat-req>1</memstat-req>
                <memstat-size>65536</memstat-size>
                <memstat-name>cluster_save buffer</memstat-name>
                <inuse>0</inuse>
                <memuse>0</memuse>
                <high-use>-</high-use>
                <memstat-req>2</memstat-req>
                <memstat-size>32,64</memstat-size>
                <memstat-name>VFS hash</memstat-name>
                <inuse>1</inuse>
                <memuse>128</memuse>
                <high-use>-</high-use>
                <memstat-req>1</memstat-req>
                <memstat-size>32,64</memstat-size>
                <memstat-name>vnodes</memstat-name>
                <inuse>1</inuse>
                <memuse>1</memuse>
                <high-use>-</high-use>
                <memstat-req>1</memstat-req>
                <memstat-size>512</memstat-size>
                <memstat-name>mount</memstat-name>
                <inuse>290</inuse>
                <memuse>23</memuse>
                <high-use>-</high-use>
                <memstat-req>535</memstat-req>
                <memstat-size>16,32,64,128,256,4096,32768</memstat-size>
                <memstat-name>vnodemarker</memstat-name>
                <inuse>0</inuse>
                <memuse>0</memuse>
                <high-use>-</high-use>
                <memstat-req>498</memstat-req>
                <memstat-size>16384</memstat-size>
                <memstat-name>pfs_nodes</memstat-name>
                <inuse>25</inuse>
                <memuse>3</memuse>
                <high-use>-</high-use>
                <memstat-req>25</memstat-req>
                <memstat-size>128</memstat-size>
                <memstat-name>pfs_vncache</memstat-name>
                <inuse>27</inuse>
                <memuse>1</memuse>
                <high-use>-</high-use>
                <memstat-req>53</memstat-req>
                <memstat-size>32</memstat-size>
                <memstat-name>STP</memstat-name>
                <inuse>1</inuse>
                <memuse>1</memuse>
                <high-use>-</high-use>
```

```
                <memstat-req>1</memstat-req>
                <memstat-size>64</memstat-size>
                <memstat-name>GEOM</memstat-name>
                <inuse>146</inuse>
                <memuse>11</memuse>
                <high-use>-</high-use>
                <memstat-req>1042</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
                <memstat-name>syncache</memstat-name>
                <inuse>1</inuse>
                <memuse>8</memuse>
                <high-use>-</high-use>
                <memstat-req>1</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
                <memstat-name>tlv_stat</memstat-name>
                <inuse>0</inuse>
                <memuse>0</memuse>
                <high-use>-</high-use>
                <memstat-req>8</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
                <memstat-name>NFS daemon</memstat-name>
                <inuse>1</inuse>
                <memuse>8</memuse>
                <high-use>-</high-use>
                <memstat-req>1</memstat-req>

<memstat-size>16,32,64,128,256,512,2048,16384,32768,131072</memstat-size>
                <memstat-name>p1003.1b</memstat-name>
                <inuse>1</inuse>
                <memuse>1</memuse>
                <high-use>-</high-use>
                <memstat-req>1</memstat-req>
                <memstat-size>16</memstat-size>
                <memstat-name>MD disk</memstat-name>
                <inuse>10</inuse>
                <memuse>20</memuse>
                <high-use>-</high-use>
                <memstat-req>10</memstat-req>
                <memstat-size>65536</memstat-size>
                <memstat-name>ata_generic</memstat-name>
                <inuse>1</inuse>
                <memuse>1</memuse>
                <high-use>-</high-use>
                <memstat-req>6</memstat-req>
                <memstat-size>16,16384,32768</memstat-size>
                <memstat-name>ISOFS mount</memstat-name>
                <inuse>8</inuse>
                <memuse>1</memuse>
                <high-use>-</high-use>
                <memstat-req>15</memstat-req>
                <memstat-size>512</memstat-size>
                <memstat-name>ISOFS node</memstat-name>
                <inuse>1440</inuse>
                <memuse>135</memuse>
                <high-use>-</high-use>
                <memstat-req>1457</memstat-req>
                <memstat-size>128</memstat-size>
                <memstat-name>CAM SIM</memstat-name>
```

```
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>-</high-use>
<memstat-req>1</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>CAM XPT</memstat-name>
<inuse>6</inuse>
<memuse>1</memuse>
<high-use>-</high-use>
<memstat-req>9</memstat-req>
<memstat-size>16,64,16384</memstat-size>
<memstat-name>CAM periph</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>-</high-use>
<memstat-req>1</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>ad_driver</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>-</high-use>
<memstat-req>1</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>pagedep</memstat-name>
<inuse>1</inuse>
<memuse>32</memuse>
<high-use>-</high-use>
<memstat-req>106</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>inodedep</memstat-name>
<inuse>1</inuse>
<memuse>128</memuse>
<high-use>-</high-use>
<memstat-req>464</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>newblk</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>-</high-use>
<memstat-req>336</memstat-req>
<memstat-size>64,4096</memstat-size>
<memstat-name>bmsafemap</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>-</high-use>
<memstat-req>63</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>allocdirect</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>-</high-use>
<memstat-req>320</memstat-req>
<memstat-size>128</memstat-size>
<memstat-name>indirdep</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>-</high-use>
<memstat-req>17</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>allocindir</memstat-name>
<inuse>0</inuse>
```

```
<memuse>0</memuse>
<high-use>-</high-use>
<memstat-req>15</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>freefrag</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>-</high-use>
<memstat-req>12</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>freeblks</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>-</high-use>
<memstat-req>40</memstat-req>
<memstat-size>2048</memstat-size>
<memstat-name>freefile</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>-</high-use>
<memstat-req>101</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>diradd</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>-</high-use>
<memstat-req>465</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>mkdir</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>-</high-use>
<memstat-req>136</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>dirrem</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>-</high-use>
<memstat-req>168</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>newdirblk</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>-</high-use>
<memstat-req>1</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>savedino</memstat-name>
<inuse>0</inuse>
<memuse>0</memuse>
<high-use>-</high-use>
<memstat-req>157</memstat-req>
<memstat-size>512</memstat-size>
<memstat-name>UFS mount</memstat-name>
<inuse>15</inuse>
<memuse>36</memuse>
<high-use>-</high-use>
<memstat-req>15</memstat-req>
<memstat-size>2048,65536,131072</memstat-size>
<memstat-name>ata_dma</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
```

```
<high-use>-</high-use>
<memstat-req>2</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>UMAHash</memstat-name>
<inuse>1</inuse>
<memuse>2</memuse>
<high-use>-</high-use>
<memstat-req>4</memstat-req>
<memstat-size>4096,16384,32768,65536</memstat-size>
<memstat-name>cdev</memstat-name>
<inuse>22</inuse>
<memuse>3</memuse>
<high-use>-</high-use>
<memstat-req>22</memstat-req>
<memstat-size>256</memstat-size>
<memstat-name>file desc</memstat-name>
<inuse>141</inuse>
<memuse>32</memuse>
<high-use>-</high-use>
<memstat-req>1583</memstat-req>
<memstat-size>16,1024,2048,16384</memstat-size>
<memstat-name>VM pgdata</memstat-name>
<inuse>2</inuse>
<memuse>65</memuse>
<high-use>-</high-use>
<memstat-req>2</memstat-req>
<memstat-size>64</memstat-size>
<memstat-name>sigio</memstat-name>
<inuse>1</inuse>
<memuse>1</memuse>
<high-use>-</high-use>
<memstat-req>20</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>kenv</memstat-name>
<inuse>24</inuse>
<memuse>5</memuse>
<high-use>-</high-use>
<memstat-req>27</memstat-req>
<memstat-size>16,32,64,131072</memstat-size>
<memstat-name>atkbddev</memstat-name>
<inuse>2</inuse>
<memuse>1</memuse>
<high-use>-</high-use>
<memstat-req>2</memstat-req>
<memstat-size>32</memstat-size>
<memstat-name>kqueue</memstat-name>
<inuse>15</inuse>
<memuse>9</memuse>
<high-use>-</high-use>
<memstat-req>19</memstat-req>
<memstat-size>1024,4096,32768</memstat-size>
<memstat-name>proc-args</memstat-name>
<inuse>57</inuse>
<memuse>3</memuse>
<high-use>-</high-use>
<memstat-req>1001</memstat-req>
<memstat-size>16,32,64,128,256,512,1024</memstat-size>
<memstat-name>isadev</memstat-name>
<inuse>21</inuse>
<memuse>2</memuse>
<high-use>-</high-use>
```

```
            <memstat-req>21</memstat-req>
            <memstat-size>64</memstat-size>
            <memstat-name>zombie</memstat-name>
            <inuse>0</inuse>
            <memuse>0</memuse>
            <high-use>-</high-use>
            <memstat-req>1278</memstat-req>
            <memstat-size>128</memstat-size>
            <memstat-name>ithread</memstat-name>
            <inuse>69</inuse>
            <memuse>5</memuse>
            <high-use>-</high-use>
            <memstat-req>69</memstat-req>
            <memstat-size>16,64,256</memstat-size>
            <memstat-name>legacydrv</memstat-name>
            <inuse>4</inuse>
            <memuse>1</memuse>
            <high-use>-</high-use>
            <memstat-req>4</memstat-req>
            <memstat-size>16</memstat-size>
            <memstat-name>memdesc</memstat-name>
            <inuse>1</inuse>
            <memuse>4</memuse>
            <high-use>-</high-use>
            <memstat-req>1</memstat-req>
            <memstat-size>131072</memstat-size>
            <memstat-name>nexusdev</memstat-name>
            <inuse>2</inuse>
            <memuse>1</memuse>
            <high-use>-</high-use>
            <memstat-req>2</memstat-req>
            <memstat-size>16</memstat-size>
            <memstat-name>CAM queue</memstat-name>
            <inuse>3</inuse>
            <memuse>1</memuse>
            <high-use>-</high-use>
            <memstat-req>3</memstat-req>
            <memstat-size>16</memstat-size>
            <memstat-name>$PIR</memstat-name>
            <inuse>4</inuse>
            <memuse>1</memuse>
            <high-use>-</high-use>
            <memstat-req>4</memstat-req>
            <memstat-size>32</memstat-size>
            <memstat-name>KTRACE</memstat-name>
            <inuse>100</inuse>
            <memuse>10</memuse>
            <high-use>-</high-use>
            <memstat-req>100</memstat-req>
            <memstat-size>128</memstat-size>
            <memstat-name>kbdmux</memstat-name>
            <inuse>5</inuse>
            <memuse>9</memuse>
            <high-use>-</high-use>
            <memstat-req>5</memstat-req>
            <memstat-size>128,2048,65536,131072</memstat-size>
        </vmstat-memstat-malloc>
        <vmstat-memstat-zone>
            <zone-name>UMA Kegs:</zone-name>
            <zone-size>136</zone-size>
            <count-limit>0</count-limit>
```

```
<used>71</used>
<free>1</free>
<zone-req>71</zone-req>
<zone-name>UMA Zones:</zone-name>
<zone-size>120</zone-size>
<count-limit>0</count-limit>
<used>71</used>
<free>19</free>
<zone-req>71</zone-req>
<zone-name>UMA Slabs:</zone-name>
<zone-size>64</zone-size>
<count-limit>0</count-limit>
<used>490</used>
<free>41</free>
<zone-req>579</zone-req>
<zone-name>UMA RCntSlabs:</zone-name>
<zone-size>104</zone-size>
<count-limit>0</count-limit>
<used>276</used>
<free>20</free>
<zone-req>276</zone-req>
<zone-name>UMA Hash:</zone-name>
<zone-size>128</zone-size>
<count-limit>0</count-limit>
<used>4</used>
<free>26</free>
<zone-req>5</zone-req>
<zone-name>16 Bucket:</zone-name>
<zone-size>76</zone-size>
<count-limit>0</count-limit>
<used>30</used>
<free>20</free>
<zone-req>30</zone-req>
<zone-name>32 Bucket:</zone-name>
<zone-size>140</zone-size>
<count-limit>0</count-limit>
<used>33</used>
<free>23</free>
<zone-req>33</zone-req>
<zone-name>64 Bucket:</zone-name>
<zone-size>268</zone-size>
<count-limit>0</count-limit>
<used>33</used>
<free>9</free>
<zone-req>33</zone-req>
<zone-name>128 Bucket:</zone-name>
<zone-size>524</zone-size>
<count-limit>0</count-limit>
<used>49</used>
<free>0</free>
<zone-req>49</zone-req>
<zone-name>VM OBJECT:</zone-name>
<zone-size>128</zone-size>
<count-limit>0</count-limit>
<used>2111</used>
<free>79</free>
<zone-req>25214</zone-req>
<zone-name>MAP:</zone-name>
<zone-size>160</zone-size>
<count-limit>0</count-limit>
<used>7</used>
```

```
<free>41</free>
<zone-req>7</zone-req>
<zone-name>KMAP ENTRY:</zone-name>
<zone-size>68</zone-size>
<count-limit>35336</count-limit>
<used>19</used>
<free>149</free>
<zone-req>2397</zone-req>
<zone-name>MAP ENTRY:</zone-name>
<zone-size>68</zone-size>
<count-limit>0</count-limit>
<used>2031</used>
<free>153</free>
<zone-req>62417</zone-req>
<zone-name>PV ENTRY:</zone-name>
<zone-size>24</zone-size>
<count-limit>509095</count-limit>
<used>57177</used>
<free>6333</free>
<zone-req>1033683</zone-req>
<zone-name>DP fakepg:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>mt_zone:</zone-name>
<zone-size>64</zone-size>
<count-limit>0</count-limit>
<used>238</used>
<free>57</free>
<zone-req>238</zone-req>
<zone-name>16:</zone-name>
<zone-size>16</zone-size>
<count-limit>0</count-limit>
<used>2114</used>
<free>119</free>
<zone-req>80515</zone-req>
<zone-name>32:</zone-name>
<zone-size>32</zone-size>
<count-limit>0</count-limit>
<used>1335</used>
<free>134</free>
<zone-req>10259</zone-req>
<zone-name>64:</zone-name>
<zone-size>64</zone-size>
<count-limit>0</count-limit>
<used>3529</used>
<free>129</free>
<zone-req>29110</zone-req>
<zone-name>96:</zone-name>
<zone-size>96</zone-size>
<count-limit>0</count-limit>
<used>2062</used>
<free>58</free>
<zone-req>4365</zone-req>
<zone-name>112:</zone-name>
<zone-size>112</zone-size>
<count-limit>0</count-limit>
<used>361</used>
<free>164</free>
```

```
<zone-req>24613</zone-req>
<zone-name>128:</zone-name>
<zone-size>128</zone-size>
<count-limit>0</count-limit>
<used>359</used>
<free>61</free>
<zone-req>942</zone-req>
<zone-name>160:</zone-name>
<zone-size>160</zone-size>
<count-limit>0</count-limit>
<used>364</used>
<free>44</free>
<zone-req>577</zone-req>
<zone-name>224:</zone-name>
<zone-size>224</zone-size>
<count-limit>0</count-limit>
<used>422</used>
<free>20</free>
<zone-req>1950</zone-req>
<zone-name>256:</zone-name>
<zone-size>256</zone-size>
<count-limit>0</count-limit>
<used>204</used>
<free>36</free>
<zone-req>1225</zone-req>
<zone-name>288:</zone-name>
<zone-size>288</zone-size>
<count-limit>0</count-limit>
<used>2</used>
<free>24</free>
<zone-req>10</zone-req>
<zone-name>512:</zone-name>
<zone-size>512</zone-size>
<count-limit>0</count-limit>
<used>49</used>
<free>7</free>
<zone-req>911</zone-req>
<zone-name>1024:</zone-name>
<zone-size>1024</zone-size>
<count-limit>0</count-limit>
<used>213</used>
<free>11</free>
<zone-req>1076</zone-req>
<zone-name>2048:</zone-name>
<zone-size>2048</zone-size>
<count-limit>0</count-limit>
<used>199</used>
<free>113</free>
<zone-req>640</zone-req>
<zone-name>4096:</zone-name>
<zone-size>4096</zone-size>
<count-limit>0</count-limit>
<used>144</used>
<free>7</free>
<zone-req>2249</zone-req>
<zone-name>Files:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>665</used>
<free>77</free>
<zone-req>16457</zone-req>
```

```
<zone-name>MAC labels:</zone-name>
<zone-size>20</zone-size>
<count-limit>0</count-limit>
<used>3998</used>
<free>227</free>
<zone-req>21947</zone-req>
<zone-name>PROC:</zone-name>
<zone-size>544</zone-size>
<count-limit>0</count-limit>
<used>116</used>
<free>10</free>
<zone-req>1394</zone-req>
<zone-name>THREAD:</zone-name>
<zone-size>416</zone-size>
<count-limit>0</count-limit>
<used>127</used>
<free>17</free>
<zone-req>131</zone-req>
<zone-name>KSEGRP:</zone-name>
<zone-size>88</zone-size>
<count-limit>0</count-limit>
<used>127</used>
<free>73</free>
<zone-req>131</zone-req>
<zone-name>UPCALL:</zone-name>
<zone-size>44</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>SLEEPQUEUE:</zone-name>
<zone-size>32</zone-size>
<count-limit>0</count-limit>
<used>145</used>
<free>194</free>
<zone-req>145</zone-req>
<zone-name>VMSPACE:</zone-name>
<zone-size>268</zone-size>
<count-limit>0</count-limit>
<used>57</used>
<free>13</free>
<zone-req>1335</zone-req>
<zone-name>mbuf_packet:</zone-name>
<zone-size>256</zone-size>
<count-limit>180000</count-limit>
<used>256</used>
<free>128</free>
<zone-req>49791</zone-req>
<zone-name>mbuf:</zone-name>
<zone-size>256</zone-size>
<count-limit>180000</count-limit>
<used>50</used>
<free>466</free>
<zone-req>105183</zone-req>
<zone-name>mbuf_cluster:</zone-name>
<zone-size>2048</zone-size>
<count-limit>25190</count-limit>
<used>387</used>
<free>165</free>
<zone-req>5976</zone-req>
<zone-name>mbuf_jumbo_pagesize:</zone-name>
```

```
<zone-size>4096</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>mbuf_jumbo_9k:</zone-name>
<zone-size>9216</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>mbuf_jumbo_16k:</zone-name>
<zone-size>16384</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>ACL UMA zone:</zone-name>
<zone-size>388</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>g_bio:</zone-name>
<zone-size>132</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>174</free>
<zone-req>69750</zone-req>
<zone-name>ata_request:</zone-name>
<zone-size>200</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>57</free>
<zone-req>5030</zone-req>
<zone-name>ata_composite:</zone-name>
<zone-size>192</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>GENCFG:</zone-name>
<zone-size>72</zone-size>
<count-limit>1000004</count-limit>
<used>57</used>
<free>102</free>
<zone-req>57</zone-req>
<zone-name>VNODE:</zone-name>
<zone-size>292</zone-size>
<count-limit>0</count-limit>
<used>2718</used>
<free>25</free>
<zone-req>2922</zone-req>
<zone-name>VNODEPOLL:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>S VFS Cache:</zone-name>
<zone-size>68</zone-size>
```

```
<count-limit>0</count-limit>
<used>2500</used>
<free>76</free>
<zone-req>3824</zone-req>
<zone-name>L VFS Cache:</zone-name>
<zone-size>291</zone-size>
<count-limit>0</count-limit>
<used>51</used>
<free>14</free>
<zone-req>63</zone-req>
<zone-name>NAMEI:</zone-name>
<zone-size>1024</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>8</free>
<zone-req>53330</zone-req>
<zone-name>NFSMOUNT:</zone-name>
<zone-size>480</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>NFSNODE:</zone-name>
<zone-size>460</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>PIPE:</zone-name>
<zone-size>404</zone-size>
<count-limit>0</count-limit>
<used>27</used>
<free>9</free>
<zone-req>717</zone-req>
<zone-name>KNOTE:</zone-name>
<zone-size>72</zone-size>
<count-limit>0</count-limit>
<used>42</used>
<free>64</free>
<zone-req>3311</zone-req>
<zone-name>socket:</zone-name>
<zone-size>412</zone-size>
<count-limit>25191</count-limit>
<used>343</used>
<free>8</free>
<zone-req>2524</zone-req>
<zone-name>unpcb:</zone-name>
<zone-size>140</zone-size>
<count-limit>25200</count-limit>
<used>170</used>
<free>26</free>
<zone-req>2157</zone-req>
<zone-name>ipq:</zone-name>
<zone-size>52</zone-size>
<count-limit>216</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>udpcb:</zone-name>
<zone-size>232</zone-size>
<count-limit>25194</count-limit>
```

```
<used>19</used>
<free>32</free>
<zone-req>31</zone-req>
<zone-name>inpcb:</zone-name>
<zone-size>232</zone-size>
<count-limit>25194</count-limit>
<used>40</used>
<free>28</free>
<zone-req>105</zone-req>
<zone-name>tcpcb:</zone-name>
<zone-size>520</zone-size>
<count-limit>25193</count-limit>
<used>40</used>
<free>16</free>
<zone-req>105</zone-req>
<zone-name>tcptw:</zone-name>
<zone-size>56</zone-size>
<count-limit>5092</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>syncache:</zone-name>
<zone-size>128</zone-size>
<count-limit>15360</count-limit>
<used>0</used>
<free>60</free>
<zone-req>55</zone-req>
<zone-name>tcpreass:</zone-name>
<zone-size>20</zone-size>
<count-limit>1690</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>sackhole:</zone-name>
<zone-size>20</zone-size>
<count-limit>0</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>ripcb:</zone-name>
<zone-size>232</zone-size>
<count-limit>25194</count-limit>
<used>5</used>
<free>29</free>
<zone-req>5</zone-req>
<zone-name>SWAPMETA:</zone-name>
<zone-size>276</zone-size>
<count-limit>94948</count-limit>
<used>0</used>
<free>0</free>
<zone-req>0</zone-req>
<zone-name>FFS inode:</zone-name>
<zone-size>132</zone-size>
<count-limit>0</count-limit>
<used>1146</used>
<free>72</free>
<zone-req>1306</zone-req>
<zone-name>FFS1 dinode:</zone-name>
<zone-size>128</zone-size>
<count-limit>0</count-limit>
<used>1146</used>
```

```
                    <free>24</free>
                    <zone-req>1306</zone-req>
                    <zone-name>FFS2 dinode:</zone-name>
                    <zone-size>256</zone-size>
                    <count-limit>0</count-limit>
                    <used>0</used>
                    <free>0</free>
                    <zone-req>0</zone-req>
                </vmstat-memstat-zone>
                <vmstat-sumstat>
                    <cpu-context-switch>934906</cpu-context-switch>
                    <dev-intr>1707986</dev-intr>
                    <soft-intr>33819</soft-intr>
                    <traps>203604</traps>
                    <sys-calls>1200636</sys-calls>
                    <kernel-thrds>60</kernel-thrds>
                    <fork-calls>1313</fork-calls>
                    <vfork-calls>21</vfork-calls>
                    <rfork-calls>0</rfork-calls>
                    <swap-pageins>0</swap-pageins>
                    <swap-pagedin>0</swap-pagedin>
                    <swap-pageouts>0</swap-pageouts>
                    <swap-pagedout>0</swap-pagedout>
                    <vnode-pageins>23094</vnode-pageins>
                    <vnode-pagedin>23119</vnode-pagedin>
                    <vnode-pageouts>226</vnode-pageouts>
                    <vnode-pagedout>3143</vnode-pagedout>
                    <page-daemon-wakeup>0</page-daemon-wakeup>
                    <page-daemon-examined-pages>0</page-daemon-examined-pages>
                    <pages-reactivated>8821</pages-reactivated>
                    <copy-on-write-faults>48364</copy-on-write-faults>
                    <copy-on-write-optimized-faults>31</copy-on-write-optimized-faults>
                    <zero-fill-pages-zeroed>74665</zero-fill-pages-zeroed>
                    <zero-fill-pages-prezeroed>70061</zero-fill-pages-prezeroed>
                    <transit-blocking-page-faults>85</transit-blocking-page-faults>
                    <total-vm-faults>191824</total-vm-faults>

<pages-affected-by-kernel-thrd-creat>0</pages-affected-by-kernel-thrd-creat>
                    <pages-affected-by-fork>95343</pages-affected-by-fork>
                    <pages-affected-by-vfork>3526</pages-affected-by-vfork>
                    <pages-affected-by-rfork>0</pages-affected-by-rfork>
                    <pages-freed>221502</pages-freed>
                    <pages-freed-by-deamon>0</pages-freed-by-deamon>
                    <pages-freed-by-exiting-proc>75630</pages-freed-by-exiting-proc>
                    <pages-active>45826</pages-active>
                    <pages-inactive>13227</pages-inactive>
                    <pages-in-vm-cache>49278</pages-in-vm-cache>
                    <pages-wired-down>10640</pages-wired-down>
                    <pages-free>70706</pages-free>
                    <bytes-per-page>4096</bytes-per-page>
                    <swap-pages-used>0</swap-pages-used>
                    <peak-swap-pages-used>0</peak-swap-pages-used>
                    <total-name-lookups>214496</total-name-lookups>
                    <positive-cache-hits>92</positive-cache-hits>
                    <negative-cache-hits>5</negative-cache-hits>
                    <pass2>0</pass2>
                    <cache-deletions>0</cache-deletions>
                    <cache-falsehits>0</cache-falsehits>
                    <toolong>0</toolong>
                </vmstat-sumstat>
                <vmstat-intr>
```

```
                <intr-name>irq0: clk          </intr-name>
                <intr-cnt>1243455</intr-cnt>
                <intr-rate>999</intr-rate>
                <intr-name>irq4: sio0         </intr-name>
                <intr-cnt>1140</intr-cnt>
                <intr-rate>0</intr-rate>
                <intr-name>irq8: rtc          </intr-name>
                <intr-cnt>159164</intr-cnt>
                <intr-rate>127</intr-rate>
                <intr-name>irq9: cbb1 fxp0    </intr-name>
                <intr-cnt>28490</intr-cnt>
                <intr-rate>22</intr-rate>
                <intr-name>irq10: fxp1        </intr-name>
                <intr-cnt>20593</intr-cnt>
                <intr-rate>16</intr-rate>
                <intr-name>irq14: ata0        </intr-name>
                <intr-cnt>5031</intr-cnt>
                <intr-rate>4</intr-rate>
                <intr-name>Total</intr-name>
                <intr-cnt>1457873</intr-cnt>
                <intr-rate>1171</intr-rate>
            </vmstat-intr>
            <vm-kernel-state>
                <vm-kmem-map-free>248524800</vm-kmem-map-free>
            </vm-kernel-state>
        </system-virtual-memory-information>
        <cli>
            <banner></banner>
        </cli>
    </rpc-reply>
```

show task replication

Syntax **show task replication**

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Displays graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) status. When you issue this command on the master Routing Engine, the status of nonstop active routing synchronization is also displayed.

Options This command has no options.

Required Privilege Level view

List of Sample Output
show task replication (Issued on the Master Routing Engine) on page 295
show task replication (Issued on the Backup Routing Engine) on page 295

Output Fields Table 42 on page 295 lists the output fields for the **show task replication** command. Output fields are listed in the approximate order in which they appear.

Table 42: show task replication Output Fields

Field Name	Field Description
Stateful replication	Displays whether or not graceful Routing Engine switchover is configured. The status can be **Enabled** or **Disabled**.
RE mode	Displays the Routing Engine on which the command is issued: **Master**, **Backup**, or **Not applicable** (when the router has only one Routing Engine).
Protocol	Protocol that are supported by nonstop active routing.
Synchronization Status	Nonstop active routing synchronization status for the supported protocols. States are **NotStarted**, **InProgress**, and **Complete**.

show task replication (Issued on the Master Routing Engine)

```
user@host> show task replication
        Stateful Replication: Enabled
        RE mode: Master

    Protocol                Synchronization Status
    OSPF                    NotStarted
    BGP                     Complete
    IS-IS                   NotStarted
    LDP                     Complete
```

show task replication (Issued on the Backup Routing Engine)

```
user@host> show task replication
        Stateful Replication: Enabled
        RE mode: Master
```

show version

Syntax
show version
<brief | detail>

Syntax (J-EX Series Switch)
show version
<all-members>
<brief | detail>
<local>
<member *member-id*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display the hostname and version information about the software running on the router or switch.

Options
none—Display standard information about the hostname and version of the software running on the router or switch.

brief | detail—(Optional) Display the specified level of output.

all-members—(J-EX4200 switches only) (Optional) Display standard information about the hostname and version of the software running on all members of the Virtual Chassis configuration.

local—(J-EX4200 switches only) (Optional) Display standard information about the hostname and version of the software running on the local Virtual Chassis member.

member *member-id*—(J-EX4200 switches only) (Optional) Display standard information about the hostname and version of the software running on the specified member of the Virtual Chassis configuration. Replace ***member-id*** with a value from 0 through 9.

Required Privilege Level
view

List of Sample Output
show version (J-EX8208) on page 297
show version (J-EX4200-24F) on page 297

show version (J-EX8208)

```
user@switch> show version
Model: DELL J-EX8208
JUNOS Base OS boot [10.2I20100329_2206_pappavu]
JUNOS Base OS Software Suite [10.2I20100329_2206_pappavu]
JUNOS Kernel Software Suite [10.2I20100329_2206_pappavu]
JUNOS Crypto Software Suite [10.2I20100329_2206_pappavu]
JUNOS Online Documentation [10.2I20100329_2206_pappavu]
JUNOS Enterprise Software Suite [10.2I20100329_2206_pappavu]
LC JUNOS Installation Software [10.2I20100329_2206_pappavu]
JUNOS Routing Software Suite [10.2I20100329_2206_pappavu]
JUNOS Web Management [10.2I20100329_2206_pappavu]
```

show version (J-EX4200-24F)

```
{master:0}

user@switch> show version
fpc0:
--------------------------------------------------------------------------
Model: DELL J-EX4200-24F
JUNOS Base OS boot [10.2I20100331_2208_pappavu]
JUNOS Base OS Software Suite [10.2I20100331_2208_pappavu]
JUNOS Kernel Software Suite [10.2I20100331_2208_pappavu]
JUNOS Crypto Software Suite [10.2I20100331_2208_pappavu]
JUNOS Online Documentation [10.2I20100331_2208_pappavu]
JUNOS Enterprise Software Suite [10.2I20100331_2208_pappavu]
JUNOS Packet Forwarding Engine Enterprise Software Suite
[10.2I20100331_2208_pappavu]
JUNOS Routing Software Suite [10.2I20100331_2208_pappavu]
JUNOS Web Management [10.2I20100331_2208_pappavu]
```

PART 6

Junos OS for J-EX Series Switches Power Management

CHAPTER 18

Power Management Overview

Junos OS—Overview

J-EX Series Switch Software Features Overview

The following tables list the J-EX Series Switches software features and the Junos OS Release in which they were introduced:

Related Documentation

- J-EX4200 Switches Hardware Overview on page 25
- J-EX8208 Switch Hardware Overview on page 27
- J-EX8216 Switch Hardware Overview on page 30

Power Management

- Understanding Power Management on J-EX Series Switches on page 302

Understanding Power Management on J-EX Series Switches

The power management feature for Dell PowerConnect J-Series J-EX8200 Ethernet Switches helps ensure that normal operation of the system is not disrupted because of insufficient power to the switch. It does so by employing a power budget policy.

Power management does the following in its power budget policy:

- Budgets power for each switch component that requires power. The amount that power management budgets for each component is the maximum power that component might consume. For example, for the fan tray, power management budgets the amount of power required to run the fans at their maximum speed setting.
- Reserves a set amount of power for power supply redundancy. In its default configuration, power management manages the switch for N+1 power redundancy, which ensures uninterrupted system operation if one power supply fails. For example, if a switch has four online 3000 W power supplies, power management reserves 3000 W in its power budget policy for redundancy. It allocates the remaining 9000 W to normal operating power.
- Specifies the rules under which components receive power. These rules are designed to ensure the least disruption to switch operation under conditions of insufficient power. For example, power management provides power to core system components, such as the Routing Engine, before it provides power to line cards.

You can configure certain aspects of power management's budget policy, specifically:

- The power priority of individual line cards. By assigning different power priorities to the line cards, you can determine which line cards are more likely to receive power in the event of insufficient power.
- The power redundancy configuration. The default power redundancy configuration is N+1; you can optionally configure N+N. For example, if you have deployed two independent AC power feeds to the switch, configure N+N redundancy. When you configure power management for N+N redundancy, it reserves the appropriate amount of power in its power budget and reports insufficient power conditions accordingly.

These configurable items are discussed further in:

- Power Priority of Line Cards on page 303
- Power Supply Redundancy on page 304

Power Priority of Line Cards

Power management powers line cards on or off based on the power priority of the slots they occupy:

- When a switch powers on, power management provides power to the line cards in the order of their slot priority, with line cards in high priority slots receiving power first. Thus if available power (including redundant power) is exhausted before all line cards receive power, higher priority cards are powered on while lower priority cards remain powered off.
- If the switch starts receiving insufficient power because of power supply failure or some other problem, power management powers off the line cards in reverse-priority order until power (including redundant power) is sufficient for the remaining cards. Thus higher priority line cards are more likely to retain power in power shortage conditions than are lower priority line cards.
- Power management responds to changes in power availability and line card operating status by powering line cards on or off as appropriate. For example, if you add a power supply, lower priority cards that were powered off due to insufficient power are powered on in priority order.

  If you take a line card offline, power management no longer allocates power to it. If power to the switch is insufficient when you take a line card offline, power management allocates the freed power to a lower priority card that was offline due to lack of power and brings it online. Restarting a line card, however, does not affect the power allocated to it. Thus when power is insufficient, restarting a line card does not change its operating status or the operating status of other line cards.

NOTE: Because power management does not allocate power to an offline line card, a line card that has been taken offline in a J-EX8200 switch is not automatically brought online when you commit a configuration. You must explicitly bring the line card online with the **request chassis fpc slot** ***slot-number*** **online** command. This behavior differs from other platforms running the Junos OS, which automatically bring an offline FPC online when you commit a configuration.

The actual power priority of a slot is determined first by the slot's assigned priority and second by the slot's number. If more than one slot has the same assigned priority, power priority is determined by slot number, with the lowest-numbered slots receiving power first.

By default, all slots are assigned the lowest priority. You can assign a priority to a slot using the CLI. If you do not explicitly assign priorities to slots, the slots receive power in ascending order of slot numbers.

Because the purpose of power management is to ensure minimal system disruption when power is insufficient, slot power priority does not always determine which line cards receive power. In some cases, power management might provide power to a lower priority line card rather than a higher priority line card. For example:

- If power is insufficient for a line card in a higher priority slot but is sufficient for a line card in a lower priority slot, the lower priority slot receives the power. For example, if an 8-port SFP+ line card requiring 450 W is in a higher priority slot than a 48-port SFP line card requiring 330 W, the 48-port SFP line card receives the power if there is more than 330 W but less than 450 W available.
- In an operating switch that has insufficient power, power management does not power off operating line cards to provide power to a newly inserted line card or a line card that is brought online after being offline, even if the line card has a higher priority than the currently operating line cards.

  However, if you restart the switch, power management reruns the current power budget policy and powers line cards on or off based on their priority. As a result, line cards receive power strictly by priority order and previously operating line cards might no longer receive power.
- If you change the assigned power priority of line cards when there is insufficient power for all the line cards, power management does not power down line cards that had been receiving power because they are now a lower priority.

Power Supply Redundancy

By default, power management in J-EX8200 switches is configured to manage the power supplies for N+1 redundancy, in which one power supply is held in reserve for backup if one of the other power supplies is removed or fails.

You can configure power management to manage the power supplies for N+N redundancy. In N+N redundancy, power management holds N power supplies in reserve for backup. For example, if your switch has six power supplies and you configure N+N redundancy, power management makes three power supplies available for normal operating power and reserves three power supplies for redundancy (3+3). If you have an odd number of power supplies, power management allocates one more power supply to normal operating power than to redundant power. For example, if you have five power supplies, the N+N configuration is 3+2.

Given the same number of power supplies, an N+N configuration usually provides less normal operating power than an N+1 configuration because the N+N configuration holds more power in reserve for backup. Table 43 on page 304 shows the effect on normal operating power in N+1 and N+N configurations.

Table 43: Available Operating Power in N+1 and N+N Redundancy Configurations

Number of Power Supplies at *n* W Each	Normal Operating Power in N+1 Configuration	Normal Operating Power in N+N Configuration
2	1 x (*n* W)	1 x (*n* W)

Table 43: Available Operating Power in N+1 and N+N Redundancy Configurations *(continued)*

Number of Power Supplies at *n* W Each	Normal Operating Power in N+1 Configuration	Normal Operating Power in N+N Configuration
3	2 x (*n* W)	2 x (*n* W)
4	3 x (*n* W)	2 x (*n* W)
5	4 x (*n* W)	3 x (*n* W)
6	5 x (*n* W)	3 x (*n* W)

To compensate for the reduced normal operating power, power management reserves less power to the chassis in an N+N configuration than in an N+1 configuration. This reduction in reserved chassis power allows a switch in an N+N configuration to power more line cards than it could without the reduction. For the J-EX8208 switch, the power reserved for the chassis is reduced to 1200 W from 1600 W; for the J-EX8216 switch, it is reduced to 1800 W from 2400 W.

NOTE: To achieve the reduction in reserved chassis power, power management reduces the maximum fan speed to 60 percent in an N+N configuration from 80 percent in an N+1 configuration. Because the maximum fan speed is reduced, it is possible that a line card that overheats would be shut down sooner in an N+N configuration than in an N+1 configuration.

Power management automatically recalculates the redundant power and normal operating power as power supplies go online or offline. For example, if you have an N+N configuration with three online 2000 W power supplies, power management allocates 2000 W to redundant power. If you bring a fourth 2000 W power supply online, power management then allocates 4000 W to redundant power. If a power supply goes offline again, power management once again allocates 2000 W to redundant power.

When power is insufficient to meet the budgeted power requirements, power management raises alarms as follows:

- If all the line cards are receiving power but insufficient redundant power exists to maintain the configured N+1 or N+N power configuration, power management raises a minor (yellow) alarm. If this condition persists for 5 minutes, the alarm becomes a major (red) alarm.
- If one or more line cards are down because of insufficient power (including redundant power), power management raises a major (red) alarm.

Power management clears all alarms when sufficient power is available to meet normal operating and redundant power requirements.

Related Documentation

- Understanding Alarm Types and Severity Levels on J-EX Series Switches on page 533

CHAPTER 19

Initial Configuration

Configuring Power Supply Redundancy (CLI Procedure)

By default, the power management feature in J-EX8200 switches is configured to manage the power supplies for N+1 redundancy, in which one power supply is held in reserve for backup if any one of the other power supplies is removed or fails.

You can configure power management to manage the power supplies for N+N redundancy. For example, to set up your AC power supplies for dual power feed, N+N redundancy is required. In N+N redundancy, power management allocates half of the online power supplies to normal operating power and half to redundant power. If you have an odd number of online power supplies, power management allocates one more power supply to normal operating power than to redundant power.

This topic describes how to configure power management for N+N redundancy and how to revert back to N+1 redundancy if your deployment needs change.

Before you configure power management for N+N redundancy, ensure that you have sufficient power supplies to meet the power requirements of this configuration. Use the **show chassis power-budget-statistics** command to display your current power budget.

NOTE: To allow more power to be available to line cards, power management compensates for the reduced normal operating power in an N+N configuration by reserving less power to the chassis than it does in an N+1 configuration. For the J-EX8208 switch, the power reserved for the chassis is reduced to 1200 W from 1600 W. For the J-EX8216 switch, it is reduced to 1800 W from 2400 W. In determining whether you have enough power for an N+N configuration, take this reduction of reserved chassis power into account.

The reduction in reserved chassis power is achieved by reducing the maximum fan speed to 60 percent in an N+N configuration from 80 percent in an N+1 configuration. Because the maximum fan speed is reduced, it is possible that a line card that overheats would be shut down sooner in an N+N configuration than in an N+1 configuration.

To configure N+N redundancy:

```
[edit chassis]
user@switch# set psu redundancy n-plus-n
```

To revert back to N+1 redundancy:

```
[edit chassis]
user@switch# delete chassis psu redundancy n-plus-n
```

Related Documentation

- Configuring the Power Priority of Line Cards (CLI Procedure) on page 308
- Verifying Power Configuration and Use on page 309
- Understanding Power Management on J-EX Series Switches on page 302

Configuring the Power Priority of Line Cards (CLI Procedure)

The power management facility on J-EX8200 switches allows you to assign power priorities to the slots occupied by line cards. Power management provides power to the slots in priority order, which means that line cards in higher priority slots are more likely to receive power than line cards in lower priority slots if power to the switch is insufficient to power all the line cards.

When assigning power priority to slots, keep these points in mind:

- 0 is the highest priority. For a J-EX8208 switch, you can assign a priority of 0 through 7 to a slot. For a J-EX8216 switch, you can assign a priority of 0 through 15 to a slot.
- All slots are assigned the lowest priority by default.
- If a group of slots shares the same assigned priority, each slot's power priority within the group is based on its slot number, with the lowest-numbered slots receiving power first.

To assign or change the power priority for a slot:

```
[edit chassis]
user@switch# set fpc slot power-budget-priority priority
```

For example, to set slot 6 to priority 0, enter:

```
[edit chassis]
user@switch# set fpc 6 power-budget-priority 0
```

Related Documentation

- Configuring Power Supply Redundancy (CLI Procedure) on page 307
- Verifying Power Configuration and Use on page 309
- Understanding Power Management on J-EX Series Switches on page 302

CHAPTER 20

Verifying Power Management

Verifying Power Configuration and Use

Purpose Verify on a J-EX8200 switch:

- What the power redundancy and line card priority settings are
- Whether the N+1 or N+N power requirements are being met
- Whether the switch has sufficient power for a new line card or an N+N configuration

Action Enter the following command:

```
user@switch> show chassis power-budget-statistics
    PSU  1     (EX8200-AC2K)                 :    1200 W
    PSU  2     (EX8200-AC2K)                 :    1200 W
    PSU  3     (EX8200-AC2K)                 :    1200 W
    PSU  4     (EX8200-AC2K)                 :    1200 W
    Total Power supplied by all Online PSUs  :    4800 W
    Power Redundancy Configuration           :     N+N
    Power Reserved for the Chassis           :    1200 W
    FPC  5     (EX8200-48F )                 :     330 W    Priority: 7
    FPC  6     (EX8200-8XS )                 :     450 W    Priority: 0
    Actual Power Used                        :    1980 W
    Power Available (Redundant case)         :     420 W
    Total Power Available                    :    2820 W
```

Meaning The switch is configured for N+N redundancy. As shown by the **Power Available (Redundant case)** field, the switch has sufficient power to meet the N+N power requirements and has an additional 420 W available. The switch has insufficient power for an additional 8-port SFP+ line card while maintaining N+N redundancy, because the line card requires 450 W. However, it does have enough power for an additional 48-port SFP line card, which requires only 330 W. The 8-port SFP+ line card in slot 6 has a higher power priority than the 48-port SFP line card line card in slot 5.

NOTE: The amount of power shown in the **Actual Power Used** field reflects the total power allocated in the power budget for the installed components rather than the actual power being used by the components. Because the power budget allocation is based on maximum power use, actual power consumption is likely to be much less.

Related Documentation

- Configuring Power Supply Redundancy (CLI Procedure) on page 307
- Configuring the Power Priority of Line Cards (CLI Procedure) on page 308

CHAPTER 21

Configuration Statements for Power Management

fpc

Syntax

```
fpc slot {
  pic pic-number {
    sfpplus {
      pic-mode mode;
    }
  }
  power-budget-priority priority;
```

Hierarchy Level [edit chassis]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description On a J-EX4200 switch, specify the port of the SFP+ uplink module for which you want to configure the operating mode.

On a J-EX8200 switch, specify the line card slot for which you want to assign a power priority.

Options ***slot***—Number of the slot:

- 0—Standalone J-EX4200 switches. The FPC refers to the switch itself.
- 0–9—J-EX4200 switch in a Virtual Chassis configuration. The value corresponds to the switch's member ID.
- 0–7—J-EX8208 switch. The slot is a line card slot.
- 0–15—J-EX8216 switch. The slot is a line card slot.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Setting the Mode on an SFP+ Uplink Module (CLI Procedure) on page 921
- Configuring the Power Priority of Line Cards (CLI Procedure) on page 308

n-plus-n

Syntax n-plus-n;

Hierarchy Level [edit **chassis psu redundancy**]

Release Information Statement introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Configure N+N power supply redundancy for power management on a J-EX8200 switch.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Power Supply Redundancy (CLI Procedure) on page 307

power-budget-priority

Syntax power-budget-priority *priority*;

Hierarchy Level [edit **chassis fpc** *slot*]

Release Information Statement introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Assign a power priority to the specified line card slot on a J-EX8200 switch.

Default All line card slots are initially assigned the lowest priority.

Options ***priority***—Assigned power priority for the slot, with 0 being the highest priority.
Range: 0 through 7 for a J-EX8208 switch; 0 through 15 for a J-EX8216 switch

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring the Power Priority of Line Cards (CLI Procedure) on page 308

psu

Syntax

```
psu {
  redundancy {
    n-plus-n;
  }
}
```

Hierarchy Level [edit chassis]

Release Information Statement introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Configure N+N power supply redundancy for power management on a J-EX8200 switch.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Power Supply Redundancy (CLI Procedure) on page 307

redundancy

Syntax

```
redundancy {
  n-plus-n;
}
```

Hierarchy Level [edit chassis psu]

Release Information Statement introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Configure N+N power supply redundancy for power management on a J-EX8200 switch.

The remaining statement is explained separately.

Default N+1 power supply redundancy is configured by default.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Power Supply Redundancy (CLI Procedure) on page 307

CHAPTER 22

Operational Mode Commands for Power Management

show chassis power-budget-statistics

Syntax **show chassis power-budget-statistics**

Release Information Command introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Display the power budget of a J-EX8200 switch.

Required Privilege Level view

Related Documentation
- Verifying Power Configuration and Use on page 309
- Configuring the Power Priority of Line Cards (CLI Procedure) on page 308
- Configuring Power Supply Redundancy (CLI Procedure) on page 307

List of Sample Output **show chassis power-budget-statistics on page 317**

Output Fields Table 44 on page 316 lists the output fields for the **show chassis power-budget-statistics** command. Output fields are listed in the approximate order in which they appear.

Table 44: show chassis power-budget-statistics Output Fields

Field Name	Field Description
PSU *n* (*supply type*)	Number of watts supplied by the power supply. All installed power supplies, whether currently operating or not, are listed.
Power supplied by all Online PSUs	Total number of watts supplied by all currently operating power supplies.
Power Redundancy Configuration	Configured power redundancy setting, either N+1 or N+N.
Power Reserved for the Chassis	Power reserved for the chassis: • For a J-EX8208 switch: 1600 W in an N+1 configuration; 1200 W in an N+N configuration • For a J-EX8216 switch: 2400 W in an N+1 configuration; 1800 W in an N+N configuration The power reserved for the chassis includes the maximum power requirements for the fan tray and Switch Fabric and Routing Engine (SRE), Routing Engine (RE), and Switch Fabric (SF) modules in both base and redundant configurations.
FPC *n* (*card type*)	Number of watts required by the line card in slot *n* and the power priority assigned to the slot.
Actual Power Used	Power budgeted for all the components in the switch. This equal to the power reserved for the chassis plus the power requirements of all online line cards. Because the amount budgeted is based on maximum power requirements, the real power consumption of the switch is likely to be less than this figure.

Table 44: show chassis power-budget-statistics Output Fields *(continued)*

Field Name	Field Description
Power Available (Redundant case)	Unused power available to the switch in the power budget, excluding redundant power. If power is insufficient to meet the N+1 or N+N redundancy requirements, this value is 0.
Total Power Available	Unused power available to the switch in the power budget, including redundant power.

show chassis power-budget-statistics

```
user@switch> show chassis power-budget-statistics
    PSU  0     (EX8200-AC2K)                 :    2000 W
    PSU  1     (EX8200-AC2K)                 :    2000 W
    PSU  2     (EX8200-AC2K)                 :    2000 W
    Total Power supplied by all Online PSUs  :    6000 W
    Power Redundancy Configuration           :     N+N
    Power Reserved for the Chassis           :    1600 W
    FPC  6     (EX8200-8XS )                 :     450 W    Priority: 7
    Actual Power Used                        :    2050 W
    Power Available (Redundant case)         :    1950 W
    Total Power Available                    :    3950 W
```

PART 7

Junos OS for J-EX Series Switches Configuration Management

CHAPTER 23

Configuration Management Overview

Configuration Files—Overview

Understanding Configuration Files for J-EX Series Switches

A configuration file stores the complete configuration of a switch. The current configuration of a switch is called the active configuration. You can alter this current configuration and you can also return to a previous configuration or to a rescue configuration. For more information, see “Configuration Files Terms” on page 322.

Junos OS saves the 50 most recently committed configuration files on the switch so that you can return to a previous configuration. The configuration files are named:

- **juniper.conf.gz**—The current active configuration.
- **juniper.conf.1.gz** to **juniper.conf.49.gz**—Rollback configurations.

To make changes to the configuration file, you have to work in the configuration mode in the CLI or use the configuration tools in the J-Web interface. When making changes to a configuration file, you are viewing and changing the candidate configuration file. The candidate configuration allows you to make configuration changes without causing operational changes to the active configuration or causing potential damage to your current network operations. Once you commit the changes made to the candidate configuration, the system updates the active configuration.

Related Documentation

- Managing Configuration Files Through the Configuration History (J-Web Procedure) on page 338
- Uploading a Configuration File (CLI Procedure) on page 336
- Uploading a Configuration File (J-Web Procedure) on page 337

- Loading a Previous Configuration File (CLI Procedure) on page 340
- Reverting to the Rescue Configuration for the J-EX Series Switch on page 343
- Configuration Files Terms on page 322

Configuration Files Terms

Table 45 on page 322 lists the various configuration file terms used for J-EX Series switches and their definitions.

Table 45: Configuration File Terms

Term	Definition
active configuration	The current committed configuration of a switch.
candidate configuration	A working copy of the configuration that allows users to make configurational changes without causing any operational changes until this copy is committed.
configuration group	Group of configuration statements that can be inherited by the rest of the configuration.
commit a configuration	Have the candidate configuration checked for proper syntax, activated, and marked as the current configuration file running on the switching platform.
configuration hierarchy	The Junos OS configuration consists of a hierarchy of statements. There are two types of statements: container statements, which contain other statements, and leaf statements, which do not contain other statements. All the container and leaf statements together form the configuration hierarchy.
default configuration	The default configuration contains the initial values set for each configuration parameter when a switch is shipped.
rescue configuration	Well-known configuration that recovers a switch from a configuration that denies management access. You set a current committed configuration to be the rescue configuration through the J-Web interface or CLI.
roll back a configuration	Return to a previously committed configuration.

Related Documentation

- J-EX4200 Default Configuration on page 325
- J-EX8200 Switch Default Configuration on page 329
- Loading a Previous Configuration File (CLI Procedure) on page 340
- Managing Configuration Files Through the Configuration History (J-Web Procedure) on page 338
- Reverting to the Rescue Configuration for the J-EX Series Switch on page 343
- Understanding Configuration Files for J-EX Series Switches on page 321

Understanding Automatic Refreshing of Scripts on J-EX Series Switches

You can automatically refresh **commit**, **event**, and **op** scripts using operational mode commands on J-EX Series switches. The commands are:

- **request system scripts refresh-from commit**
- **request system scripts refresh-from event**
- **request system scripts refresh-from op**

The existing Junos OS command-line interface (CLI) **refresh** and **refresh-from** configuration mode statements have been extended to work with Junos XML management protocol and NETCONF XML management protocol sessions.

Related Documentation

- Understanding Autoinstallation of Configuration Files on J-EX Series Switches on page 323
- CLI User Interface Overview on page 127
- *Junos OS Junos XML Management Protocol Guide* at http://www.juniper.net/techpubs/software/junos/
- *Junos OS NETCONF XML Management Protocol Guide*at http://www.juniper.net/techpubs/software/junos/

Understanding Autoinstallation of Configuration Files on J-EX Series Switches

Autoinstallation is the automatic configuration of a device over the network from a pre-existing configuration file that you create and store on a configuration server—typically a Trivial File Transfer Protocol (TFTP) server. You can use autoinstallation to automatically configure new devices and to deploy multiple devices from a central location in the network.

Autoinstallation takes place automatically when you connect an Ethernet port on a new switch to the network and power on the switch. You can also explicitly enable autoinstallation on J-EX Series Switches in your network to implement autoinstallation when they are powered on. To configure autointallation, you specify a configuration server, an autoinstallation interface, and a protocol for IP address acquisition.

This topic describes:

- Typical Uses for Autoinstallation on page 323
- Autoinstallation Configuration Files and IP Addresses on page 324
- Typical Autoinstallation Process on a New Switch on page 324

Typical Uses for Autoinstallation

- To deploy and update multiple devices from a central location in the network.
- To configure a new device—Autoinstallation takes place when you power on a device that has only the factory default configuration (boot) file.

- To update a device—Autoinstallation takes place when a device that has been manually configured for autoinstallation is powered on.

Autoinstallation Configuration Files and IP Addresses

For the autoinstallation process to work, you must store one or more host-specific or default configuration files on a configuration server in the network and have a service available—typically Dynamic Host Configuration Protocol (DHCP)—to assign an IP address to the switch.

You can set up the following configuration files for autoinstallation on the switch:

- **network.conf**—Default configuration file for autoinstallation, in which you specify IP addresses and associated hostnames for devices on the network.
- **switch.conf**—Default configuration file for autoinstallation with a minimum configuration sufficient for you to telnet to the device and configure it manually.
- ***hostname*.conf**—Host-specific configuration file for autoinstallation on a device that contains all the configuration information necessary for the switch. In the filename, ***hostname*** is replaced with the hostname assigned to the switch.

If the server with the autoinstallation configuration file is not on the same LAN segment as the new device, or if a specific device is required by the network, you must configure an intermediate device directly attached to the new switch, through which the new switch can send TFTP, boot protocol (BOOTP), and Domain Name System (DNS) requests. In this case, you specify the IP address of the intermediate device as the location to receive TFTP requests for autoinstallation.

Typical Autoinstallation Process on a New Switch

When a J-EX Series switch is powered on for the first time, it performs the following autoinstallation tasks:

1. The new switch sends out DHCP or BOOTP requests on each connected interface simultaneously to obtain an IP address.

   If a DHCP server responds to these requests, it provides the switch with some or all of the following information:

   - An IP address and subnet mask for the autoinstallation interface.
   - The location of the (typically) TFTP server, Hypertext Transfer Protocol (HTTP) server, or FTP server on which the configuration file is stored.
   - The name of the configuration file to be requested from the TFTP server.
   - The IP address or hostname of the TFTP server.

     If the DHCP server provides the server's hostname, a DNS server must be available on the network to resolve the name to an IP address.
   - The IP address of an intermediate device if the configuration server is on a different LAN segment from the new switch.

2. After the new switch acquires an IP address, the autoinstallation process on the switch attempts to download a configuration file in the following ways:

   a. If the DHCP server specifies the host-specific configuration file ***hostname*.conf**, the switch uses that filename in the TFTP server request. The autoinstallation process on the new switch makes three unicast TFTP requests for ***hostname*.conf**. If these attempts fail, the switch broadcasts three requests to any available TFTP server for the file.

   b. If the new switch does not locate a ***hostname*.conf** file, the autoinstallation process sends three unicast TFTP requests for a **network.conf** file that contains the switch's hostname-to-IP-address mapping information. If these attempts fail, the switch broadcasts three requests to any available TFTP server for the file.

   c. If the switch fails to find a **network.conf** file that contains a hostname entry for the switch, the autoinstallation process sends out a DNS request and attempts to resolve the new switch's IP address to a hostname.

   d. If the new switch determines its hostname, it sends a TFTP request for the ***hostname*.conf** file.

   e. If the new switch is unable to map its IP address to a hostname, it sends TFTP requests for the default configuration file **switch.conf**. The TFTP request procedure is the same as for the **network.conf** file.

3. After the new switch locates a configuration file on a TFTP server, the autoinstallation process downloads the file, installs the file on the switch, and commits the configuration.

Related Documentation

- Configuring Autoinstallation of Configuration Files (CLI Procedure) on page 345
- Connecting and Configuring a J-EX Series Switch (CLI Procedure) on page 161
- Connecting and Configuring a J-EX Series Switch (J-Web Procedure) on page 163
- Configuration Files Terms on page 322

J-EX Series Switches Default Configuration

- J-EX4200 Default Configuration on page 325
- J-EX8200 Switch Default Configuration on page 329

J-EX4200 Default Configuration

Each J-EX Series switch is programmed with a factory default configuration that contains the values set for each configuration parameter when a switch is shipped. The default configuration file sets values for system parameters such as **syslog** and **commit**; configures Power over Ethernet (PoE), storm control, and Ethernet switching on all interfaces; and enables the LLDP and RSTP protocols.

When you commit changes to the configuration, a new configuration file is created that becomes the active configuration. You can always revert to the factory default

configuration. See “Reverting to the Default Factory Configuration for the J-EX Series Switch” on page 341.

The following factory default configuration file is for a J-EX4200 switch with 24 ports (for models that have more ports, this default configuration file has more interfaces):

NOTE: In this example, ge-0/0/0 through ge-0/0/23 are the network interface ports. Optional uplink modules provide four 1-gigabit SFP transceivers (ge-0/1/0 through ge-0/1/3). Although you can install only one uplink module, the interfaces for both are shown below.

```
system {
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
    commit {
        factory-settings {
            reset-chassis-lcd-menu;
            reset-virtual-chassis-configuration;
        }
    }
}
interfaces {
    ge-0/0/0 {
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/0/1 {
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/0/2 {
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/0/3 {
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/0/4 {
```

```
    unit 0 {
        family ethernet-switching;
    }
}
ge-0/0/5 {
    unit 0 {
        family ethernet-switching;
    }
}
ge-0/0/6 {
    unit 0 {
        family ethernet-switching;
    }
}
ge-0/0/7 {
    unit 0 {
        family ethernet-switching;
    }
}
ge-0/0/8 {
    unit 0 {
        family ethernet-switching;
    }
}
ge-0/0/9 {
    unit 0 {
        family ethernet-switching;
    }
}
ge-0/0/10 {
    unit 0 {
        family ethernet-switching;
    }
}
ge-0/0/11 {
    unit 0 {
        family ethernet-switching;
    }
}
ge-0/0/12 {
    unit 0 {
        family ethernet-switching;
    }
}
ge-0/0/13 {
    unit 0 {
        family ethernet-switching;
    }
}
ge-0/0/14 {
    unit 0 {
        family ethernet-switching;
    }
}
ge-0/0/15 {
    unit 0 {
```

```
        family ethernet-switching;
    }
}
ge-0/0/16 {
    unit 0 {
        family ethernet-switching;
    }
}
ge-0/0/17 {
    unit 0 {
        family ethernet-switching;
    }
}
ge-0/0/18 {
    unit 0 {
        family ethernet-switching;
    }
}
ge-0/0/19 {
    unit 0 {
        family ethernet-switching;
    }
}
ge-0/0/20 {
    unit 0 {
        family ethernet-switching;
    }
}
ge-0/0/21 {
    unit 0 {
        family ethernet-switching;
    }
}
ge-0/0/22 {
    unit 0 {
        family ethernet-switching;
    }
}
ge-0/0/23 {
    unit 0 {
        family ethernet-switching;
    }
}
ge-0/1/0 {
    unit 0 {
        family ethernet-switching;
    }
}
ge-0/1/1 {
    unit 0 {
        family ethernet-switching;
    }
}
ge-0/1/2 {
    unit 0 {
        family ethernet-switching;
```

```
        }
      }
      ge-0/1/3 {
        unit 0 {
          family ethernet-switching;
        }
      }
    }
    protocols {
      igmp-snooping{
        vlan all;
      }
      lldp {
        interface all;
      }
      lldp-med {
        interface all;
      }
      rstp;
    }
    ethernet-switching-options {
      storm-control {
        interface all;
      }
    }
    poe {
      interface all;
    }
```

Related Documentation

- Reverting to the Default Factory Configuration for the J-EX Series Switch on page 341
- Connecting and Configuring a J-EX Series Switch (CLI Procedure) on page 161
- Connecting and Configuring a J-EX Series Switch (J-Web Procedure) on page 163
- Understanding Configuration Files for J-EX Series Switches on page 321
- J-EX Series Switches Interfaces Overview on page 863

J-EX8200 Switch Default Configuration

Each J-EX8200 switch is programmed with a factory default configuration that contains the values set for each configuration parameter when a switch is shipped. The default configuration file sets values for system parameters such as the ARP aging timer, the system log, and file messages, while also enabling the LLDP protocol, the RSTP protocol, IGMP snooping, and storm control.

When you commit changes to the configuration, a new configuration file is created that becomes the active configuration. You can always revert to the factory default configuration. See “Reverting to the Default Factory Configuration for the J-EX Series Switch” on page 341.

This topic shows the factory default configuration file of a J-EX8200 switch:

```
system {
  arp {
```

```
        aging-timer 5
        }
    }
    syslog {
        user * {
            any emergency;
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
    commit {
        factory-settings {
            reset-chassis-lcd-menu;
        }
    }
}
    protocols {
        igmp-snooping {
            vlan all;
        }
        lldp {
            interface all;
        }
        rstp;
    }
    ethernet-switching-options {
        storm-control {
            interface all;
        }
    }
```

Related Documentation

- Configuration Files Terms on page 322
- Connecting and Configuring a J-EX Series Switch (CLI Procedure) on page 161
- Connecting and Configuring a J-EX Series Switch (J-Web Procedure) on page 163
- Understanding Configuration Files for J-EX Series Switches on page 321
- J-EX8208 Switch Hardware Overview on page 27
- J-EX8216 Switch Hardware Overview on page 30

CHAPTER 24

Managing Junos OS Configuration

Using the Configuration Tools in J-Web

Using the CLI Viewer in the J-Web Interface to View Configuration Text

To view the entire configuration file contents in text format, select **Configure>CLI Tools >CLI Viewer**. The main pane displays the configuration in text format.

Each level in the hierarchy is indented to indicate each statement's relative position in the hierarchy. Each level is generally set off with braces, with an open brace ({) at the beginning of each hierarchy level and a closing brace (}) at the end. If the statement at a hierarchy level is empty, the braces are not displayed. Each leaf statement ends with a semicolon (;), as does the last statement in the hierarchy.

This indented representation is used when the configuration is displayed or saved as an ASCII file. However, when you load an ASCII configuration file, the format of the file is not so strict. The braces and semicolons are required, but the indention and use of new lines are not required in ASCII configuration files.

Related Documentation

- Understanding J-Web Configuration Tools on page 131

Using the CLI Editor in the J-Web Interface to Edit Configuration Text

Use the CLI Editor to edit configuration if you know the Junos OS CLI or prefer a command interface.

To edit the entire configuration in text format:

CAUTION: We recommend that you use this method to edit and commit the configuration only if you have experience editing configurations through the CLI.

1. Select **Configure>CLI Tools>CLI Editor**. The main pane displays the configuration in a text editor.
2. Navigate to the hierarchy level you want to edit.

   You can edit the candidate configuration using standard text editor operations—insert lines (by using the Enter key), delete lines, and modify, copy, and paste text.
3. Click **Commit** to load and commit the configuration.

   The switching platform checks the configuration for the correct syntax before committing it.

Related Documentation

- CLI User Interface Overview on page 127
- Understanding J-Web Configuration Tools on page 131

Using the Point and Click CLI Tool in the J-Web Interface to Edit Configuration Text

To edit the configuration on a series of pages of clickable options that steps you through the hierarchy, select **Configure>CLI Tools>Point&Click CLI**. The side pane displays the top level of the configured hierarchy, and the main pane displays configured hierarchy options and the Icon Legend.

To expand or hide the hierarchy of all the statements in the side pane, click **Expand all** or **Hide all**. To expand or hide an individual statement in the hierarchy, click the expand (+) or collapse (–) icon to the left of the statement.

TIP: Only those statements included in the committed configuration are displayed in the hierarchy.

The configuration information in the main pane consists of configuration options that correspond to configuration statements. Configuration options that contain subordinate statements are identified by the term *Nested*.

To include, edit, or delete statements in the candidate configuration, click one of the links described in Table 46 on page 332. Then specify configuration information by typing in a field, selecting a value from a list, or clicking a check box (toggle).

Table 46: J-Web Edit Point & Click Configuration Links

Link	Function
Add new entry	Displays fields and lists for a statement identifier, allowing you to add a new identifier to a statement.

Table 46: J-Web Edit Point & Click Configuration Links *(continued)*

Link	Function
Configure	Displays information for a configuration option that has not been configured, allowing you to include a statement.
Delete	Deletes the corresponding statement or identifier from the configuration. All subordinate statements and identifiers contained within a deleted statement are also discarded.
Edit	Displays information for a configuration option that has already been configured, allowing you to edit a statement.
Identifier	Displays fields and lists for an existing statement identifier, allowing you to edit the identifier.

As you navigate through the configuration, the hierarchy level is displayed at the top of the main pane. You can click a statement or identifier in the hierarchy to display the corresponding configuration options in the main pane.

The main pane includes icons that display information about statements and identifiers when you place your cursor over them. Table 47 on page 333 describes these icons.

Table 47: J-Web Edit Point & Click Configuration Icons

Icon	Function
C	Displays a comment about a statement.
I	Indicates that a statement is inactive.
M	Indicates that a statement has been added or modified but has not been committed.
*	Indicates that the statement or identifier is required in the configuration.
?	Provides online help information.

After typing or selecting your configuration edits, click a button in the main pane (described in Table 48 on page 333) to apply your changes or cancel them, refresh the display, or discard parts of the candidate configuration. An updated configuration does not take effect until you commit it.

Table 48: J-Web Edit Point & Click Configuration Buttons

Button	Function
Refresh	Updates the display with any changes to the configuration made by other users.
Commit	Verifies edits and applies them to the current configuration file running on the switch.
Discard	Removes edits applied to or deletes existing statements or identifiers from the candidate configuration.

Related Documentation

- CLI User Interface Overview on page 127
- Understanding J-Web Configuration Tools on page 131

Using the Commit Options to Commit Configuration Changes (J-Web Procedure)

You can use the single-commit feature to commit all outstanding configuration changes in the J-Web interface on J-EX Series switches simultaneously. This helps in reducing the time J-Web takes for committing configurations because when changes are committed at every step, rollback configurations pile up.

For example, suppose you want to delete a firewall filter and add a new one. With immediate commits, you would need to commit your changes twice for this action. Using single commit, you can decrease the number of commits to one, thus saving time for working on other configurations.

When you edit a configuration, you work on a copy of the current configuration, which is your candidate configuration. The changes you make to the candidate configuration are visible through the user interface immediately, allowing other users to edit those configurations, but they do not take effect on the switch until you commit the changes. When you commit the configuration, the candidate file is checked for proper syntax, activated, and marked as the current, operational software configuration file. If multiple users are editing the configuration when you commit the candidate configuration, changes made by all users take effect.

You can configure the commit options to either commit all configuration changes together or commit each configuration change immediately using the J-Web Commit Preference page.

NOTE: There are some pages on which configuration changes must be committed immediately. For such pages, if you configure the commit options for a single commit, the system displays warning notifications that remind you to commit your changes immediately. An example of such a page is the Interface Page (**Configure** > **Interface**).

To configure the commit options on a J-EX Series switch using the J-Web interface:

1. Select **Commit Options**.

   

   NOTE: All action links except **Preference** are disabled unless you edit, add, or delete a configuration.

2. Choose an action. See Table 49 on page 335 for details on the actions.
3. Configure the commit options by selecting **Preference**. See Table 50 on page 335 for details on preference options.

Table 49: Commit Options

Menu Item	Function	Your Action
Commit	Commits the candidate configuration of the current user session, along with changes from other user sessions.	1. Select **Commit Options** > **Commit**. Changes are committed after the system validates your configuration. A window displays that the configuration was successfully committed or that the commit failed. 2. Click **OK**. Click **Details** to view the commit log.
Compare	Displays the XML log of pending uncommitted configurations on the device.	1. Select **Commit Options** > **Compare**. The XML log of pending configurations on the devices are displayed similar to the CLI interface, in a "human-readable" form. 2. Click **Close**.
Discard	Discards the candidate configuration of your current session, along with changes from other user sessions.	1. Select **Commit Options** > **Discard**. 2. Click **OK** to confirm the discard action. Your changes are discarded after the system validates your configuration.
Preference	Indicates your choice of committing all global configurations together or committing each configuration change immediately.	1. Select **Commit Options** > **Preference**. The Commit Preference page is displayed. 2. Configure the commit options by selecting your preference. See Table 50 on page 335 for details on preference options.

Table 50: Commit Preference Options

Option	Function
Validate and commit configuration changes	Sets the system to validate and force an immediate commit on every screen after every configuration change.
Validate configuration changes	Loads all the configuration changes for an accumulated single commit. If there are errors in loading the configuration, the errors are logged. This is the default mode. Once you select this option, you need to select **Commit Options** > **Commit** to commit your changes.

Related Documentation

- J-Web User Interface for J-EX Series Switches Overview on page 129
- J-EX Series Switch Software Features Overview on page 3

Managing Junos OS Configuration

- Uploading a Configuration File (CLI Procedure) on page 336
- Uploading a Configuration File (J-Web Procedure) on page 337

Uploading a Configuration File (CLI Procedure)

You can create a configuration file on your local system, copy the file to the J-EX Series switch and then load the file into the CLI. After you have loaded the configuration file, you can commit it to activate the configuration on the switch. You can also edit the configuration interactively using the CLI and commit it at a later time.

To upload a configuration file from your local system:

1. Create the configuration file using a text editor such as Notepad, making sure that the syntax of the configuration file is correct. For more information about testing the syntax of a configuration file see *Junos OS System Basics and Services Command Reference* at http://www.juniper.net/techpubs/software/junos/.
2. In the configuration text file, use an option to perform the required action when the file is loaded. Table 51 on page 336 lists and describes some options for the **load** command.

Table 51: Options for the load command

Options	Description
merge	Combines the current active configuration and the configuration in *filename* or the one that you type at the terminal. A **merge** operation is useful when you are adding a new section to an existing configuration. If the active configuration and the incoming configuration contain conflicting statements, the statements in the incoming configuration override those in the active configuration.
override	Discards the current candidate configuration and loads the configuration in *filename* or the one that you type at the terminal. When you use the **override** option and commit the configuration, all system processes reparse the configuration. You can use the **override** option at any level of the hierarchy.
replace	Searches for the **replace** tags, deletes the existing statements of the same name, if any, and replaces them with the incoming configuration. If there is no existing statement of the same name, the **replace** operation adds the statements marked with the **replace** tag to the active configuration. NOTE: For this operation to work, you must include **replace** tags in the text file or in the configuration you type at the terminal.

3. Press Ctrl+A to select all the text in the configuration file.
4. Press Ctrl+C to copy the contents of the configuration text file to the Clipboard.

5. Log in to the switch using your username and password.
6. To enter configuration mode:
   user@switch> configure

   You will see this output, with the hash or pound mark indicating configuration mode.
   Entering configuration mode
   [edit]
   user@switch#
7. Load the configuration file:
   [edit]
   user@switch# load merge terminal
8. At the cursor, paste the contents of the Clipboard using the mouse and the Paste icon:
   [edit]
   user@switch# load merge terminal
   [Type ^D at a new line to end input]
   >Cursor is here. Paste the contents of the clipboard here<
9. Press Enter.
10. Press Ctrl+D to set the end-of-file marker.

To view results of the configuration steps before committing the configuration, type the **show** command at the user prompt.

To commit these changes to the active configuration, type the **commit** command at the user prompt.You can also edit the configuration interactively using the CLI and commit it at a later time.

Related Documentation

- Uploading a Configuration File (J-Web Procedure) on page 337
- Understanding Configuration Files for J-EX Series Switches on page 321

Uploading a Configuration File (J-Web Procedure)

You can create a configuration file on your local system, copy the file to the J-EX Series switch and then load the file into the CLI. After you have loaded the configuration file, you can commit it to activate the configuration on the switch. You can also edit the configuration interactively using the CLI and commit it at a later time.

To upload a configuration file from your local system:

1. Select **Maintain > Config Management > Upload**.

   The main pane displays the File to Upload box.
2. Specify the name of the file to upload using one of the following methods:
   - Type the absolute path and filename in the File to Upload box.
   - Click **Browse** to navigate to the file.
3. Click **Upload and Commit** to upload and commit the configuration.

The switch checks the configuration for the correct syntax before committing it.

Related Documentation

- Uploading a Configuration File (CLI Procedure) on page 336
- Understanding J-Web Configuration Tools on page 131
- Understanding Configuration Files for J-EX Series Switches on page 321

Managing Configuration Files Through the Configuration History (J-Web Procedure)

Use the Configuration History function to manage configuration files.

1. Displaying Configuration History on page 338
2. Displaying Users Editing the Configuration on page 339
3. Comparing Configuration Files with the J-Web Interface on page 339
4. Downloading a Configuration File with the J-Web Interface on page 340
5. Loading a Previous Configuration File with the J-Web Interface on page 340

Displaying Configuration History

To manage configuration files with the J-Web interface, select **Maintain > Config Management >History**. The main pane displays History — Database Information page.

Table 52 on page 338 summarizes the contents of the display.

The configuration history display allows you to:

- View a configuration.
- Compare two configurations.
- Download a configuration file to your local system.
- Roll back the configuration to any of the previous versions stored on the switch.

Table 52: J-Web Configuration History Summary

Field	Description
Number	Version of the configuration file.
Date/Time	Date and time the configuration was committed.
User	Name of the user who committed the configuration.
Client	Method by which the configuration was committed: • **cli**—A user entered a Junos OS CLI command. • **junoscript**—A Junos XML protocol client performed the operation. Commit operations performed by users through the J-Web interface are identified in this way. • **snmp**—An SNMP **set** request started the operation. • **other**—Another method was used to commit the configuration.
Comment	Comment.

Table 52: J-Web Configuration History Summary *(continued)*

Field	Description
Log Message	Method used to edit the configuration: • Imported via paste— Configuration was edited and loaded with the **Configure>CLI Tools>Edit Configuration Text** option. • Imported upload [*filename*]—Configuration was uploaded with the **Configure>CLI Tools>Point Click Editor** option. • Modified via J–Web Configure — Configuration was modified with the J-Web Configure menu. • Rolled back via *user-interface*— Configuration was rolled back to a previous version through the user interface specified by *user-interface*, which can be Web Interface or CLI.
Action	Action to perform with the configuration file. The action can be **Download** or **Rollback**.

Displaying Users Editing the Configuration

To display a list of users editing the switching platform configuration, select **Config Management >History**. The list is displayed as Database Information in the main pane. Table 53 on page 339 summarizes the Database Information display.

Table 53: J-Web Configuration Database Information Summary

Field	Description
User Name	Name of user editing the configuration.
Start Time	Time of day the user logged in to the switch.
Idle Time	Elapsed time since the user issued a configuration command from the CLI.
Terminal	Terminal on which the user is logged in.
PID	Process identifier assigned to the user by the switching platform.
Edit Flags	Designates a private or exclusive edit.
Edit Path	Level of the configuration hierarchy that the user is editing.

Comparing Configuration Files with the J-Web Interface

To compare any two of the past 50 committed configuration files:

1. Select **Config Management >History**. A list of the current and the previous 49 configurations is displayed as Configuration History in the main pane.
2. Select the check boxes to the left of the two configuration versions you want to compare.
3. Click **Compare**.

   The main pane displays the differences between the two configuration files at each hierarchy level as follows:

- Lines that have changed are highlighted side by side in green.
- Lines that exist only in the more recent configuration file are displayed in red on the left.
- Lines that exist only in the older configuration file are displayed in blue on the right.

Downloading a Configuration File with the J-Web Interface

To download a configuration file from the switch to your local system:

1. Select **Config Management >History**. A list of current and previous 49 configurations is displayed as Configuration History in the main pane.
2. In the Action column, click **Download** for the version of the configuration you want to download.
3. Select the options your Web browser provides that allow you to save the configuration file to a target directory on your local system.

   The file is saved as an ASCII file.

Loading a Previous Configuration File with the J-Web Interface

To load (roll back) and commit a previous configuration file stored on the switching platform:

1. Select **Config Management >History**. A list of current and previous 49 configurations is displayed as Configuration History in the main pane.
2. In the Action column, click **Rollback** for the version of the configuration you want to load.

   The main pane displays the results of the rollback operation.

NOTE: When you click Rollback, the switch loads and commits the selected configuration. This behavior is different from the switch's behavior that occurs after you enter the rollback configuration mode command from the CLI. In the latter case, the configuration is loaded but not committed.

Related Documentation

- Loading a Previous Configuration File (CLI Procedure) on page 340
- Understanding Configuration Files for J-EX Series Switches on page 321
- Understanding J-Web Configuration Tools on page 131

Loading a Previous Configuration File (CLI Procedure)

You can return to a previously committed configuration file if you need to revert to a previous configuration. The J-EX Series switch saves the last 50 committed configurations, including the rollback number, date, time, and name of the user who issued the **commit** configuration command.

Syntax

rollback <*number*>

Options

- **none**— Return to the most recently saved configuration.
- ***number***—Configuration to return to.
  - **Range**: 0 through 49. The most recently saved configuration is number 0, and the oldest saved configuration is number 49.
  - **Default**: 0

To return to a configuration prior to the most recently committed one:

1. Specify the rollback number (here, 1 is entered and the configuration returns to the previously committed configuation):

   ```
   [edit]
   user@switch# rollback 1
   load complete
   ```

2. Activate the configuration you have loaded:

   ```
   [edit]
   user@switch# commit
   ```

Related Documentation

- Managing Configuration Files Through the Configuration History (J-Web Procedure) on page 338
- Configuration Files Terms on page 322
- For more information on **rollback**, see the *Junos OS CLI User Guide* at http://www.juniper.net/techpubs/software/junos/ .

Reverting to the Default Factory Configuration for the J-EX Series Switch

If for any reason the current active configuration fails, you can revert to the default factory configuration. You can also roll back to a previous configuration, as described in "Loading a Previous Configuration File (CLI Procedure)" on page 340, or revert to the rescue configuration, as described in "Reverting to the Rescue Configuration for the J-EX Series Switch" on page 343.

The default factory configuration contains the basic configuration settings. This is the first configuration of the switch and it is loaded when the switch is first installed and powered on.

You can revert to the default factory configuration by using the **Menu** button to the right of the LCD on the front panel of the switch or by using the **load factory default** configuration command.

- Reverting to the Default Factory Configuration by Using the LCD Panel on page 342
- Reverting to the Default Factory Configuration by Using the Load Factory Default Command on page 342

Reverting to the Default Factory Configuration by Using the LCD Panel

To set the switch to the default factory configuration, use the LCD panel and buttons on the front panel of the switch as shown in Figure 5 on page 342. If the switch model does not have an LCD panel, use the CLI command described in the following section.

Use the LCD panel to revert to the default factory configuration if you want to run EZsetup. When you use the CLI to revert to the default factory configuration, the configuration for the root password is retained and you cannot run EZSetup.

Figure 5: J-EX Series Switch LCD Panel

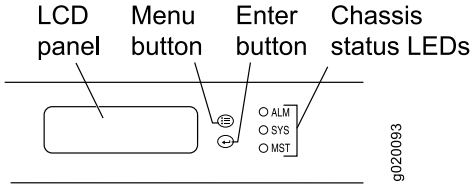

NOTE: If you want to convert a J-EX4200 switch from a member of a multimember Virtual Chassis configuration to a standalone switch, first disconnect the cables connected to the Virtual Chassis ports (VCPs). See Disconnecting a Virtual Chassis Cable from a J-EX4200 Switch. The Menu button procedure deletes all modified configuration parameters, including Virtual Chassis parameters such as member ID, mastership priority, and setting of VCP uplinks.

1. Press the **Menu** button until you see MAINTENANCE MENU on the panel.
2. Press the **Enter** button.
3. Press **Menu** until you see FACTORY DEFAULT.
4. Press **Enter**. The display says RESTORE DEFAULT?
5. Press **Enter**. The screen flashes **FACTORY DEFAULT IN PROGRESS** and returns to the idle menu.

Reverting to the Default Factory Configuration by Using the Load Factory Default Command

The **load factory default** command is a standard Junos OS configuration command. This configuration command replaces the current active configuration with the default factory configuration.

Use the LCD panel to revert to the default factory configuration if you want to run EZsetup. When you use the CLI to revert to the default factory configuration, the configuration for the root password is retained and you cannot run EZSetup.

NOTE: The **load factory default** command by itself is not supported on J-EX4200 switches configured in a Virtual Chassis with multiple members. In a multimember Virtual Chassis configuration, you can revert to the default factory configuration while retaining the Virtual Chassis parameters (member ID, mastership priority, or settings of VCP uplinks) using the following procedure:

1. [edit]
   user@switch# **load factory default**
2. [edit]
   user@switch# **delete system commit factory-settings**
3. [edit]
   user@switch# **commit**
4. Check the member ID and mastership priority with the **show virtual-chassis status** command and check to see whether there are remaining settings for uplink VCPs by using the **show virtual-chassis vc-port** command.

Related Documentation

- Configuring a Virtual Chassis (CLI Procedure) on page 781
- J-EX4200 Default Configuration on page 325
- J-EX8200 Switch Default Configuration on page 329
- Understanding Configuration Files for J-EX Series Switches on page 321
- For more information about the **load factory default** command see the *Junos OS CLI User Guide* at http://www.juniper.net/techpubs/software/junos/.

Reverting to the Rescue Configuration for the J-EX Series Switch

If someone inadvertently commits a configuration that denies management access to a J-EX Series switch and the console port is not accessible, you can overwrite the invalid configuration and replace it with the rescue configuration by using the LCD panel on the switch. The rescue configuration is a previously committed, valid configuration.

You can also revert to the default factory configuration, as described in "Reverting to the Default Factory Configuration for the J-EX Series Switch" on page 341.

Before you begin to revert to the rescue configuration:

- Ensure that you have physical access to the switch.
- A rescue configuration for the switch must have been previously set.

To revert the switch to the rescue configuration:

1. At the LCD panel on the switch, press **Menu** until you see **MAINTENANCE MENU**.
2. Press **Enter**.
3. Press **Menu** until you see **Load Rescue**.
4. Press **Enter**.
5. When **Commit Rescue** is displayed, press **Enter**.

   The LCD panel displays the message **Commit Rescue in Progress**. When the reversion is complete, it displays the idle menu.

NOTE: If there is no rescue configuration saved on the switch, the message Commit rescue failed is displayed.

Related Documentation

- Setting or Deleting the Rescue Configuration (CLI Procedure) on page 344
- Setting or Deleting the Rescue Configuration (J-Web Procedure) on page 345
- LCD Panel in J-EX4200 Switches
- LCD Panel in a J-EX8200 Switch
- Configuration Files Terms on page 322

Setting or Deleting the Rescue Configuration (CLI Procedure)

A rescue configuration is a well-known configuration that recovers a switch from a configuration that denies management access. You set a current committed configuration to be the rescue configuration through the J-Web interface or CLI.

If someone inadvertently commits a configuration that denies management access to a J-EX Series switch and the console port is not accessible, you can overwrite the invalid configuration and replace it with the rescue configuration by using the LCD panel on the switch. The rescue configuration is a previously committed, valid configuration. We recommend that the rescue configuration include the IP address (accessible from the network) for the management port.

To set the current active configuration as the rescue configuration:

```
user@switch> request system configuration rescue save
```

To delete an existing rescue configuration:

```
user@switch> request system configuration rescue delete
```

Related Documentation

- Setting or Deleting the Rescue Configuration (J-Web Procedure) on page 345
- Reverting to the Rescue Configuration for the J-EX Series Switch on page 343
- Loading a Previous Configuration File (CLI Procedure) on page 340

- Configuration Files Terms on page 322
- For information on show system configuration rescue, see the *Junos OS System Basics and Services Command Reference* at http://www.juniper.net/techpubs/software/junos/ .

Setting or Deleting the Rescue Configuration (J-Web Procedure)

A rescue configuration is a well-known configuration that recovers a switch from a configuration that denies management access. You set a current committed configuration to be the rescue configuration through the J-Web interface or CLI.

If someone inadvertently commits a configuration that denies management access to a J-EX Series switch and the console port is not accessible, you can overwrite the invalid configuration and replace it with the rescue configuration by using the LCD panel on the switch. The rescue configuration is a previously committed, valid configuration. We recommend that the rescue configuration include the IP address (accessible from the network) for the management port.

To view, set, or delete the rescue configuration using the J-Web interface, select **Maintain > Config Management >Rescue**. On the Rescue page, you can perform the following tasks:

- View the current rescue configuration—Click **View rescue configuration**.
- Set the current running configuration as the rescue configuration—Click **Set rescue configuration**.
- Delete the current rescue configuration—Click **Delete rescue configuration**.

Related Documentation

- Setting or Deleting the Rescue Configuration (CLI Procedure) on page 344
- Reverting to the Rescue Configuration for the J-EX Series Switch on page 343
- Configuration Files Terms on page 322

Configuring Autoinstallation of Configuration Files (CLI Procedure)

Autoinstallation is the automatic configuration of a device over the network from a pre-existing configuration file that you create and store on a configuration server—typically a Trivial File Transfer Protocol (TFTP) server. You can use autoinstallation to automatically configure new devices and to deploy multiple devices from a central location in the network.

No configuration is required on a new switch (a switch that has the factory default configuration file), because it is an automated process. However, to specify autoinstallation to run when you power on a switch already installed in your network, you can enable it by specifying one or more interfaces, protocols, and configuration servers to be used for autoinstallation.

Before you explicitly enable and configure autoinstallation on the switch, perform these tasks as needed for your network's configuration:

- Have a service available—typically Dynamic Host Configuration Protocol (DHCP)—to assign an IP address to the switch
- Configure a DHCP server on your network to meet your network requirements. You can configure a J-EX Series switch to operate as a DHCP server. For more information, see “Configuring DHCP Services (J-Web Procedure)” on page 447.
- Create one of the following configuration files, and store it on a TFTP server (or HTTP server or FTP server) in the network:
  - A host-specific file with the name ***hostname*.conf** for each switch undergoing autoinstallation. Replace ***hostname*** with the name of a switch. The ***hostname*.conf** file typically contains all the configuration information necessary for the switch with this hostname.
  - A default configuration file named **switch.conf** with the minimum configuration necessary to enable you to telnet into the new switch for further configuration.
- Physically attach the switch to the network using a Gigabit Ethernet port.
- If you configure the DHCP server to provide only the TFTP server hostname, add an IP address-to-hostname mapping entry for the TFTP server to the DNS database file on the Domain Name System (DNS) server in the network.
- If the new switch is not on the same network segment as the DHCP server (or other device providing IP address resolution), configure an existing device as an intermediate device to receive TFTP and DNS requests and forward them to the TFTP server and the DNS server. You must configure the LAN or serial interface on the intermediate device with the IP addresses of the hosts providing TFTP and DNS services. Connect this interface to the new switch.
- If you are using ***hostname*.conf** files for autoinstallation, you must also complete the following tasks:
  - Configure the DHCP server to provide a ***hostname*.conf** filename to each new switch. Each switch uses its ***hostname*.conf** filename to request a configuration file from the TFTP server. Copy the necessary ***hostname*.conf** configuration files to the TFTP server.
  - Create a default configuration file named **network.conf**, and copy it to the TFTP server. This file contains IP-address-to-hostname mapping entries. If the DHCP server does not send a ***hostname*.conf** filename to a new switch, the switch uses **network.conf** to resolve its hostname based on its IP address.

    Alternatively, you can add the IP-address-to-hostname mapping entry for the new switch to a DNS database file.

    The switch uses the hostname to request a ***hostname*.conf** file from the TFTP server.

To configure autoinstallation:

1. Specify the URL address of one or more servers from which to obtain configuration files.

```
[edit system]
user@switch# set autoinstallation configuration-servers tftp://tftpconfig.sp.com
```

NOTE: You can also use an FTP address, for example, ftp://*user*:*password*@sftpconfig.sp.com.

2. Configure one or more Ethernet interfaces to perform autoinstallation and one or two procurement protocols for each interface. The switch uses the protocols to send a request for an IP address for the interface:

```
[edit system]
user@switch# set autoinstallation interfaces ge-0/0/0 bootp
```

Related Documentation

- Verifying Autoinstallation Status on a J-EX Series Switch on page 349
- Understanding Autoinstallation of Configuration Files on J-EX Series Switches on page 323
- DHCP Services for J-EX Series Switches Overview on page 445

CHAPTER 25

Verifying Configuration

Verifying Autoinstallation Status on a J-EX Series Switch

Purpose Display the status of the autoinstallation feature on a J-EX Series switch.

Action From the CLI, enter the **show system autoinstallation status** command.

Sample Output

```
user@switch> show system autoinstallation status
Autoinstallation status:
Master state: Active
Last committed file: None
Configuration server of last committed file: 10.25.100.1
Interface:
  Name: ge-0/0/0
  State: Configuration Acquisition
  Acquired:
    Address: 192.168.124.75
    Hostname: host-ge-000
    Hostname source: DNS
    Configuration filename: switch-ge-000.conf
    Configuration filename server: 10.25.100.3
  Address acquisition:
    Protocol: DHCP Client
    Acquired address: None
    Protocol: RARP Client
    Acquired address: None
Interface:
  Name: ge-0/0/1
  State: None
  Address acquisition:
    Protocol: DHCP Client
    Acquired address: None
    Protocol: RARP Client
    Acquired address: None
```

Meaning The output shows the settings configured for autoinstallation. Verify that the values displayed are correct for the switch when it is deployed on the network.

Related Documentation

- Configuring Autoinstallation of Configuration Files (CLI Procedure) on page 345

CHAPTER 26

Configuration Statements for Configuration Management

archival

Syntax

```
archival {
  configuration {
    archive-sites {
      ftp://username:<password>@<host>:<port>/<url-path>;
      scp://<username>:<password>@<host>:<port>/<url-path>;
    }
    transfer-interval interval;
    transfer-on-commit;
  }
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure copying of the currently active configuration to an archive site.

Options The remaining statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Using Junos OS to Configure a Router or Switch to Transfer Its Configuration to an Archive Site

archive-sites (Configuration File)

Syntax

```
archive-sites {
    file://<path>/<filename>;
    ftp://username@host:<port>url-path password password;
    http://username@host:<port>url-path password password;
    scp://username@host:<port>url-path password password;
}
```

Hierarchy Level [edit system archival **configuration**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify where to transfer the current configuration files. When specifying a URL in a Junos OS statement using an IPv6 host address, you must enclose the entire URL in quotation marks (" ") and enclose the IPv6 host address in brackets ([ ]). For example, **"scp://*username*<:*password*>@[*ipv6-host-address*]<:*port*>/*url-path*"**

If you specify more than one archive site, the router or switch attempts to transfer the configuration files to the first archive site in the list, moving to the next only if the transfer fails. The format for the destination filename is ***router-name*_juniper.conf[.gz]_*YYYYMMDD_HHMMSS***.

NOTE: The time included in the destination filename is always in Coordinated Universal Time (UTC) regardless of whether the time on the router or switch is configured as UTC or the local time zone. The default time zone on the router or switch is UTC.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Using Junos OS to Configure a Router or Switch to Transfer Its Configuration to an Archive Site
- **configuration on page 355**
- **transfer-on-commit on page 358**

autoinstallation

Syntax

```
autoinstallation {
  configuration-servers {
    url;
  }
  interfaces {
    interface-name {
      bootp;
      rarp;
    }
  }
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For J-EX Series switches only. Download a configuration file automatically from an FTP, Hypertext Transfer Protocol (HTTP), or Trivial FTP (TFTP) server. When you power on a router or switch configured for autoinstallation, it requests an IP address from a Dynamic Host Configuration Protocol (DHCP) server. Once the router or switch has an address, it sends a request to a configuration server and downloads and installs a configuration.

Options The remaining statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Upgrading Software Using Automatic Software Download on J-EX Series Switches on page 82
- **configuration-servers on page 356**
- **idle-timeout on page 418**

commit synchronize

Syntax commit synchronize;

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For devices with multiple Routing Engines only. Configure a **commit** command to automatically result in a **commit synchronize** command. The Routing Engine on which you execute the **commit** command (the requesting Routing Engine) copies and loads its candidate configuration to the other (the responding) Routing Engines. All Routing Engines then perform a syntax check on the candidate configuration file being committed. If no errors are found, the configuration is activated and becomes the current operational configuration on all Routing Engines.

Accounting of events and operations on a backup Routing Engine is not supported on accounting servers such as TACACS+ or RADIUS. Logging of accounting events is supported only for events and operations on a master Routing Engine.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring Multiple Routing Engines to Synchronize Committed Configurations Automatically

configuration

Syntax

```
configuration {
  transfer-interval interval;
  transfer-on-commit;
  archive-sites {
    file://<path>/<filename>;
    ftp://<username>:<password>@<host>:<port>/<url-path> password password;
    http://username@host:<port>url-path password password;
    scp://username@host:<port>url-path password password;
  }
}
```

Hierarchy Level [edit system archival]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the router or switch to transfer its currently active configuration by means of FTP periodically or after each commit.

Options The remaining statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Using Junos OS to Configure a Router or Switch to Transfer Its Configuration to an Archive Site
- **archive on page 560**
- **transfer-interval on page 358**
- **transfer-on-commit on page 358**

configuration-servers

Syntax

```
configuration-servers {
  url;
}
```

Hierarchy Level [edit system **autoinstallation**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For J-EX Series switches only, configure the URL address of a server from which to obtain configuration files. Examples of URLs:

tftp://*hostname/path/filename*

ftp://*username:prompt@ftp.hostname.net/filename /*

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Upgrading Software Using Automatic Software Download on J-EX Series Switches on page 82
- **autoinstallation on page 353**
- **idle-timeout on page 418**

interfaces

Syntax

```
interfaces {
  interface-name {
    bootp;
    rarp;
    slarp;
  }
}
```

Hierarchy Level [edit system **autoinstallation**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For J-EX Series switches only. Configure the interface on which to perform autoinstallation. A request for an IP address is sent from the interface. Specify the IP address procurement protocol.

Options **bootp**—Send requests over serial interfaces with Frame Relay.

rarp—Send requests over Ethernet interfaces.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Upgrading Software Using Automatic Software Download on J-EX Series Switches on page 82
- *J Series Services Router Basic LAN and WAN Access Configuration Guide*
- **autoinstallation on page 353**

transfer-interval (Configuration)

Syntax transfer-interval *interval*;

Hierarchy Level [edit system archival configuration]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the switch to periodically transfer its currently active configuration to an archive site.

Options ***interval***—Interval at which to transfer the current configuration to an archive site.
Range: **15** through **2880** minutes

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- **archive on page 560**
- **configuration on page 355**
- **transfer-on-commit on page 358**

transfer-on-commit

Syntax transfer-on-commit;

Hierarchy Level [edit system archival configuration]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the switch to transfer its currently active configuration to an archive site each time you commit a candidate configuration.

NOTE: When specifying a URL in a Junos OS statement using an IPv6 host address, you must enclose the entire URL in quotation marks (“ ”) and enclose the IPv6 host address in brackets ([ ]). For example, “ftp://*username*<:*password*>@[*ipv6-host-address*]<:*port*>/*url-path*”

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- **archive on page 560**
- **configuration on page 355**
- **transfer-interval on page 358**

CHAPTER 27

Operational Mode Commands for Configuration Management

clear log

Syntax clear log *filename*
<all>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Remove contents of a log file.

Options *filename*—Name of the specific log file to truncate.

all—(Optional) Truncate the specified log file and delete all archived versions of it.

Required Privilege Level clear

Related Documentation
- **show log on page 657**

List of Sample Output **clear log on page 360**

Output Fields See **file list** for an explanation of output fields.

clear log The following sample commands list log file information, clear the contents of a log file, and then display the updated log file information:

```
user@host> file list lcc0-re0:/var/log/sampled detail
lcc0-re0:
--------------------------------------------------------------------------
-rw-r-----  1 root  wheel      26450 Jun 23 18:47 /var/log/sampled
total 1

user@host> clear log lcc0-re0:sampled
lcc0-re0:
--------------------------------------------------------------------------

user@host> file list lcc0-re0:/var/log/sampled detail
lcc0-re0:
--------------------------------------------------------------------------
-rw-r-----  1 root  wheel         57 Sep 15 03:44 /var/log/sampled
total 1
```

clear system commit

Syntax clear system commit

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Clear any pending commit operation.

Options This command has no options.

Required Privilege Level maintenance (or the actual user who scheduled the commit)

Related Documentation
- show system commit on page 381

List of Sample Output
clear system commit on page 361
clear system commit (None Pending) on page 361
clear system commit (User Does Not Have Required Privilege Level) on page 361

Output Fields When you enter this command, you are provided feedback on the status of your request.

clear system commit

```
user@host> clear system commit
Pending commit cleared.
```

clear system commit (None Pending)

```
user@host> clear system commit
No commit scheduled.
```

clear system commit (User Does Not Have Required Privilege Level)

```
user@host> clear system commit
error: Permission denied
```

file archive

Syntax file archive destination *destination* source *source* <compress>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Archive, and optionally compress, one or multiple local system files as a single file, locally or at a remote location.

Options

destination *destination*—Destination of the archived file or files. Specify the destination as a URL or filename. The Junos OS adds one of the following suffixes if the destination filename does not already have it:

- For archived files—The suffix **.tar**
- For archived and compressed files—The suffix **.tgz**

source *source*—Source of the original file or files. Specify the source as a URL or filename.

compress—(Optional) Compress the archived file with the GNU zip (gzip) compression utility. The compressed files have the suffix **.tgz**.

Required Privilege Level maintenance

List of Sample Output

file archive (Multiple Files) on page 362
file archive (Single File) on page 362
file archive (with Compression) on page 363

Output Fields When you enter this command, you are provided feedback on the status of your request.

file archive (Multiple Files) The following sample command archives all message files in the local directory **/var/log/messages** as the single file **messages-archive.tar** in the same directory:

```
user@host> file archive source /var/log/messages* destination /var/log/messages-archive.tar
/usr/bin/tar: Removing leading / from absolute path names in the archive.

user@host>
```

file archive (Single File) The following sample command archives one message file in the local directory **/var/log/messages** as the single file **messages-archive.tar** in the same directory:

```
user@host> file archive source /var/log/messages destination /var/log/messages-archive.tar
/usr/bin/tar: Removing leading / from absolute path names in the archive.
user@host
```

file archive (with Compression)

The following sample command archives and compresses all message files in the local directory **/var/log/messages** as the single file **messages-archive.tgz** in the same directory:

```
user@host> file archive compress source /var/log/messages* destination
/var/log/messages-archive.tgz
/usr/bin/tar: Removing leading / from absolute path names in the archive.

user@host>
```

file checksum md5

Syntax file checksum md5 <*pathname*> *filename*

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Calculate the Message Digest 5 (MD5) checksum of a file.

Options *pathname*—(Optional) Path to a filename.

filename—Name of a local file for which to calculate the MD5 checksum.

Required Privilege Level maintenance

Related Documentation

- Configuring Checksum Hashes for a Commit Script in the *Junos OS Configuration and Diagnostic Automation Guide*
- Configuring Checksum Hashes for an Event Script in the *Junos OS Configuration and Diagnostic Automation Guide*
- Configuring Checksum Hashes for an Op Script in the *Junos OS Configuration and Diagnostic Automation Guide*
- Executing an Op Script from a Remote Site in the *Junos OS Configuration and Diagnostic Automation Guide*
- **file checksum sha-256 on page 366**
- **file checksum sha1 on page 365**
- **op on page 207**

List of Sample Output **file checksum md5 on page 364**

Output Fields When you enter this command, you are provided feedback on the status of your request.

file checksum md5

```
user@host> file checksum md5 jbundle-5.3R2.4-export-signed.tgz
MD5 (jbundle-5.3R2.4-export-signed.tgz) = 2a3b69e43f9bd4893729cc16f505a0f5
```

file checksum sha1

Syntax file checksum sha1 <*pathname*> *filename*

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Calculate the Secure Hash Algorithm (SHA-1) checksum of a file.

Options *pathname*—(Optional) Path to a filename.

filename—Name of a local file for which to calculate the SHA-1 checksum.

Required Privilege Level maintenance

Related Documentation

- Configuring Checksum Hashes for a Commit Script in the *Junos OS Configuration and Diagnostic Automation Guide*
- Configuring Checksum Hashes for an Event Script in the *Junos OS Configuration and Diagnostic Automation Guide*
- Configuring Checksum Hashes for an Op Script in the *Junos OS Configuration and Diagnostic Automation Guide*
- Executing an Op Script from a Remote Site in the *Junos OS Configuration and Diagnostic Automation Guide*
- **file checksum md5 on page 364**
- **file checksum sha-256 on page 366**
- **op on page 207**

List of Sample Output **file checksum sha1 on page 365**

Output Fields When you enter this command, you are provided feedback on the status of your request.

file checksum sha1

```
user@host> file checksum sha1 /var/db/scripts/opscript.slax

SHA1 (/var/db/scripts/commitscript.slax) = ba9e47120c7ce55cff29afd73eacd370e162c676
```

file checksum sha-256

Syntax file checksum sha-256 <*pathname*> *filename*

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Calculate the Secure Hash Algorithm 2 family (SHA-256) checksum of a file.

Options *pathname*—(Optional) Path to a filename.

filename—Name of a local file for which to calculate the SHA-256 checksum.

Required Privilege Level maintenance

Related Documentation

- Configuring Checksum Hashes for a Commit Script in the *Junos OS Configuration and Diagnostic Automation Guide*
- Configuring Checksum Hashes for an Event Script in the *Junos OS Configuration and Diagnostic Automation Guide*
- Configuring Checksum Hashes for an Op Script in the *Junos OS Configuration and Diagnostic Automation Guide*
- Executing an Op Script from a Remote Site in the *Junos OS Configuration and Diagnostic Automation Guide*
- **file checksum md5 on page 364**
- **file checksum sha1 on page 365**
- **op on page 207**

List of Sample Output **file checksum sha-256 on page 366**

Output Fields When you enter this command, you are provided feedback on the status of your request.

file checksum sha-256

```
user@host> file checksum sha-256 /var/db/scripts/commitscript.slax

SHA256 (/var/db/scripts/commitscript.slax) =
94c2b061fb55399e15babd2529453815601a602b5c98e5c12ed929c9d343dd71
```

file compare

Syntax
file compare (files *filename filename*)
< context | unified>
<ignore-white-space >

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Compare two local files and describe the differences between them in default, context, or unified output styles:

- **Default**—In the first line of output, **c** means lines were changed between the two files, **d** means lines were deleted between the two files, and **a** means lines were added between the two files. The numbers preceding this alphabetical marker represent the first file, and the lines after the alphabetical marker represent the second file. A left angle bracket (**<**) in front of output lines refers to the first file. A right angle bracket (**>**) in front of output lines refers to the second file.
- **Context**—Display is divided into two parts. The first part is the first file; the second part is the second file. Output lines preceded by an exclamation point (**!**) have changed. Additions are marked with a plus sign (**+**), and deletions are marked with a minus sign (**-**).
- **Unified**—Display is preceded by the line number from the first and the second file (*xx*,*xxx*,*x*). Before the line number, additions to the file are marked with a plus sign (**+**), and deletions to the file are marked with a minus sign (**-**). The body of the output contains the affected lines. Changes are viewed as additions plus deletions.

Options
files *filename*—Names of two local files to compare.

context—(Optional) Display output in context format.

ignore-white-space—(Optional) Ignore changes in amount of white space.

unified—(Optional) Display output in unified format.

Required Privilege Level
none

List of Sample Output
file compare files on page 368
file compare files context on page 368
file compare files unified on page 368
file compare files unified ignore-white-space on page 369

Output Fields
When you enter this command, you are provided feedback on the status of your request.

file compare files

```
user@host> file compare files /tmp/one /tmp/two
100c100
<             full-name "File 1";
---
>             full-name "File 2";
102c102
<             class foo; # 'foo' is not defined
---
>             class super-user;
```

file compare files context

```
user@host> file compare files /tmp/one /tmp/two context
*** /tmp/one    Wed Dec  3 17:12:50 2003
--- /tmp/two    Wed Dec  3 09:13:14 2003
***************
*** 97,104 ****
                }
            }
            user bill {
!               full-name "Bill Smith";
!               class foo; # 'foo' is not defined
                authentication {
                    encrypted-password SECRET;
                }
--- 97,105 ----
                }
            }
            user bill {
!               full-name "Bill Smith";
!               uid 1089;
!               class super-user;
                authentication {
                    encrypted-password SECRET;
                }
```

file compare files unified

```
user@host> file compare files /tmp/one /tmp/two unified
--- /tmp/one    Wed Dec  3 17:12:50 2003
+++ /tmp/two    Wed Dec  3 09:13:14 2003
@@ -97,8 +97,9 @@
                }
            }
            user bill {
-               full-name "Bill Smith";
-               class foo; # 'foo' is not defined
+               full-name "Bill Smith";
+               uid 1089;
+               class super-user;
                authentication {
                    encrypted-passwordSECRET;
                }
```

file compare files unified ignore-white-space

```
user@host> file compare files /tmp/one /tmp/two unified ignore-white-space
--- /tmp/one    Wed Dec  3 09:13:10 2003
+++ /tmp/two    Wed Dec  3 09:13:14 2003
@@ -99,7 +99,7 @@
         user bill {
             full-name "Bill Smith";
             uid 1089;
-            class foo; # 'foo' is not defined
+            class super-user;
             authentication {
                  encrypted-password <SECRET>; # SECRET-DATA
             }
```

file copy

Syntax

file copy *source destination*
<source-address *address*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Copy files from one place to another on the local router or switch or between the local router or switch and a remote system.

Options

source—Source of the original file. Specify this as a URL or filename.

destination—Destination of the copied file. Specify this as a URL or filename. If you are copying a file to the current directory (your home directory on the local router or switch) and are not renaming the file, specify the destination with a period (.).

source-address *address*—(Optional) Source IP host address. This option is useful for specifying the source address of a secure copy (scp) file transfer.

Required Privilege Level

maintenance

List of Sample Output

file copy (A File from the Router to a PC) on page 370
file copy (A Configuration File Between Routing Engines) on page 370
file copy (A Log File Between Routing Engines) on page 370

Output Fields

When you enter this command, you are provided feedback on the status of your request.

file copy (A File from the Router to a PC)

```
user@host> file copy /var/tmp/rpd.core.4 berry:/c/junipero/tmp

...transferring.file..... |          0 KB |   0.3 kB/s | ETA: 00:00:00 | 100%
```

file copy (A Configuration File Between Routing Engines)

The following sample command copies a configuration file from Routing Engine 0 to Routing Engine 1:

```
user@host>  file copy /config/juniper.conf re1:/var/tmp/copied-juniper.conf
```

file copy (A Log File Between Routing Engines)

The following sample command copies a log file from Routing Engine 0 to Routing Engine 1:

```
user@host> file copy lcc0-re0:/var/log/chassisd lcc0-re1:/var/tmp
```

file delete

Syntax
file delete *filename*
<purge>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Delete a file on the local router or switch.

Options
filename—Name of the file to delete. For a routing matrix, include chassis information in the filename if the file to be deleted is not local to the Routing Engine from which the command is issued.

purge—(Optional) Overwrite regular files before deleting them.

Required Privilege Level
maintenance

List of Sample Output
file delete on page 371
file delete (Routing Matrix) on page 371

Output Fields
When you enter this command, you are provided feedback on the status of your request.

file delete

```
user@host> file list /var/tmp
dcd.core
rpd.core
snmpd.core

user@host> file delete /var/tmp/snmpd.core
user@host> file list /var/tmp
dcd.core
rpd.core
```

file delete (Routing Matrix)

```
user@host> file list lcc0-re0:/var/tmp
dcd.core
rpd.core
snmpd.core

user@host> file delete lcc0-re0:/var/tmp/snmpd.core
user@host> file list /var/tmp
dcd.core
rpd.core
```

file list

Syntax
file list
<detail | recursive>
<*filename*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display a list of files on the local router or switch.

Options
none—Display a list of all files for the current directory.

detail | recursive—(Optional) Display detailed output or descend recursively through the directory hierarchy, respectively.

filename—(Optional) Display a list of files. For a routing matrix, the filename must include the chassis information.

Additional Information
The default directory is the home directory of the user logged into the router or switch. To view available directories, enter a space and then a backslash (/) after the **file list** command. To view files within a specific directory, include a backslash followed by the directory and, optionally, subdirectory name after the **file list** command.

Required Privilege Level
maintenance

List of Sample Output
file list on page 372

Output Fields
When you enter this command, you are provided feedback on the status of your request.

file list

```
user@host> file list /var/tmp
dcd.core
rpd.core
snmpd.core
```

file rename

Syntax file rename *source destination*

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Rename a file on the local router or switch.

Options *destination*—New name for the file.

source—Original name of the file. For a routing matrix, the filename must include the chassis information.

Required Privilege Level maintenance

List of Sample Output **file rename on page 373**

Output Fields When you enter this command, you are provided feedback on the status of your request.

file rename The following example lists the files in **/var/tmp**, renames one of the files, and then displays the list of files again to reveal the newly named file.

```
user@host> file list /var/tmp
dcd.core
rpd.core
snmpd.core

user@host> file rename /var/tmp/dcd.core /var/tmp/dcd.core.990413
user@host> file list /var/tmp
dcd.core.990413
rpd.core
snmpd.core
```

file show

Syntax file show *filename*
<encoding base64>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display the contents of a file.

Options *filename*—Name of a file. For a routing matrix, the filename must include the chassis information.

encoding base64—(Optional) Encode file contents.

Required Privilege Level maintenance

List of Sample Output file show on page 375

Output Fields When you enter this command, you are provided feedback on the status of your request.

file show

```
user@host> file show /var/log/messages
Apr 13 21:00:08 romney /kernel: so-1/1/2: loopback suspected; going to standby.
Apr 13 21:00:40 romney /kernel: so-1/1/2: loopback suspected; going to standby.
Apr 13 21:02:48 romney last message repeated 4 times
Apr 13 21:07:04 romney last message repeated 8 times
Apr 13 21:07:13 romney /kernel: so-1/1/0: Clearing SONET alarm(s) RDI-P
Apr 13 21:07:29 romney /kernel: so-1/1/0: Asserting SONET alarm(s) RDI-P
...
```

request system configuration rescue delete

Syntax request system configuration rescue delete

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Delete an existing rescue configuration.

Options This command has no options.

Required Privilege Level maintenance

Related Documentation

- request system configuration rescue save on page 377
- request system software rollback on page 113
- show system commit on page 381

List of Sample Output request system configuration rescue delete on page 376

Output Fields This command produces no output.

request system configuration rescue delete

```
user@host> request system configuration rescue delete
```

request system configuration rescue save

Syntax request system configuration rescue save

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Save the most recently committed configuration as the rescue configuration so that you can return to it at any time by using the **rollback** command.

Options This command has no options.

Required Privilege Level maintenance

Related Documentation

- **request system software delete on page 111**
- **request system software rollback on page 113**
- **show system commit on page 381**

List of Sample Output **request system configuration rescue save on page 377**

Output Fields This command produces no output.

request system configuration rescue save

```
user@host> request system configuration rescue save
```

request system scripts refresh-from commit

Syntax **request system scripts refresh-from commit file *file-name* url *url-path***

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Automatically download the initial Junos OS configuration and a set of standard commit scripts during a Junos XML management protocol/NETCONF session when a switch is brought up for the first time.

The Junos XML management protocol equivalent for this operational mode command is:

```
<request-script-refresh-from>
   <type>commit</type>
   <file>file-name></file>
   <URL>URL</URL>
</request-script-refresh-from>
```

Options file *file-name*—Name of the file to be downloaded.

url *url-path*—URL of the file to be downloaded.

Required Privilege Level maintenance

Related Documentation

- Understanding Automatic Refreshing of Scripts on J-EX Series Switches on page 323
- *Junos OS Junos XML Management Protocol Guide* at http://www.juniper.net/techpubs/software/junos/
- *Junos OS NETCONF XML Management Protocol Guide* at http://www.juniper.net/techpubs/software/junos/

List of Sample Output **request system scripts refresh-from commit file config.txt url http://host1.juniper.net on page 378**

request system scripts refresh-from commit file config.txt url http://host1.juniper.net

```
user@switch> request system scripts refresh-from commit file config.txt url
http://host1.juniper.net
user@switch>
```

request system scripts refresh-from event

Syntax **request system scripts refresh-from event file *file-name* url *url-path***

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Automatically download the initial Junos OS configuration and a set of standard event scripts during a Junos XML management protocol/NETCONF session when a switch is brought up for the first time.

The Junos XML management protocol equivalent for this operational mode command is:

```
<request-script-refresh-from>
   <type>event</type>
   <file>file-name></file>
   <URL>URL</URL>
</request-script-refresh-from>
```

Options file *file-name*—Name of the file to be downloaded.

url *url-path*—URL of the file to be downloaded.

Required Privilege Level maintenance

Related Documentation

- Understanding Automatic Refreshing of Scripts on J-EX Series Switches on page 323
- *Junos OS Junos XML Management Protocol Guide* at http://www.juniper.net/techpubs/software/junos/
- *Junos OS NETCONF XML Management Protocol Guide* at http://www.juniper.net/techpubs/software/junos/

List of Sample Output **request system scripts refresh-from event file config.txt url http://host1.juniper.net on page 379**

request system scripts refresh-from event file config.txt url http://host1.juniper.net

```
user@switch> request system scripts refresh-from event file config.txt url http://host1.juniper.net
user@switch>
```

request system scripts refresh-from op

Syntax **request system scripts refresh-from op file *file-name* url *url-path***

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Automatically download the initial Junos OS configuration and a set of standard op scripts during a Junos XML management protocol/NETCONF session when a switch is brought up for the first time.

The Junos XML management protocol equivalent for this operational mode command is:

```
<request-script-refresh-from>
   <type>op</type>
   <file>file-name></file>
   <URL>URL</URL>
</request-script-refresh-from>
```

Options file *file-name*—Name of the file to be downloaded.

url *url-path*—URL of the file to be downloaded.

Required Privilege Level maintenance

Related Documentation

- Understanding Automatic Refreshing of Scripts on J-EX Series Switches on page 323
- *Junos OS Junos XML Management Protocol Guide* at http://www.juniper.net/techpubs/software/junos/
- *Junos OS NETCONF XML Management Protocol Guide* at http://www.juniper.net/techpubs/software/junos/

List of Sample Output **request system scripts refresh-from op file config.txt url http://host1.juniper.net on page 380**

request system scripts refresh-from op file config.txt url http://host1.juniper.net

```
user@switch> request system scripts refresh-from op file config.txt url http://host1.juniper.net
user@switch>
```

show system commit

Syntax show system commit

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display the pending commit operation (if any) and the commit history.

Options This command has no options.

Required Privilege Level view

Related Documentation
- **clear system commit on page 361**

List of Sample Output
show system commit on page 382
show system commit (At a Particular Time) on page 382
show system commit (At the Next Reboot) on page 382
show system commit (Rollback Pending) on page 382

Output Fields Table 54 on page 381 describes the output fields for the **show system commit** command. Output fields are listed in the approximate order in which they appear.

Table 54: show system commit Output Fields

Field Name	Field Description
Commit History	Displays the last 50 commit operations listed, most recent to first. The identifier **rescue** designates a configuration created for recovery using the **request system configuration rescue save** command.
Timestamp	Date and time of the commit operation.
User name	User who executed the commit operation
Commit method	Method used to execute the commit operation: • **cli**—CLI interactive user performed the commit operation. • **junoscript**—Junos XML protocol client performed the commit operation. • **synchronize**—The **commit synchronize** command was performed on the other Routing Engine. • **snmp**—An SNMP SET request caused the commit operation. • **button**—A button on the router or switch was pressed to commit a rescue configuration for recovery. • **autoinstall**—A configuration obtained through autoinstallation was committed. • **other**—A method other than those identified was used to perform the commit operation.

show system commit

```
user@host> show system commit
0   2003-07-28 19:14:04 PDT by root via other
1   2003-07-25 22:01:36 PDT by regress via cli
2   2003-07-25 22:01:32 PDT by regress via cli
3   2003-07-25 21:30:13 PDT by root via button
4   2003-07-25 13:46:48 PDT by regress via cli
5   2003-07-25 05:33:21 PDT by root via autoinstall
...
rescue  2002-05-10 15:32:03 PDT by root via other
```

show system commit (At a Particular Time)

```
user@host> show system commit
commit requested by root via cli at Tue May  7 15:59:00 2002
```

show system commit (At the Next Reboot)

```
user@host> show system commit
commit requested by root via cli at reboot
```

show system commit (Rollback Pending)

```
user@host> show system commit
0 2005-01-05 15:00:37 PST by root via cli commit confirmed, rollback in 3mins
```

show system configuration archival

Syntax show system configuration archival

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display directory and number of files queued for archival transfer.

Options This command has no options.

Required Privilege Level maintenance

List of Sample Output **show system configuration archival on page 383**

show system configuration archival

```
user@host> show system configuration archival

/var/transfer/config/:
total 8
```

show system configuration rescue

Syntax show system configuration rescue

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display a rescue configuration, if one exists.

Options This command has no options.

Required Privilege Level maintenance

List of Sample Output show system configuration rescue on page 384

show system configuration rescue

```
user@host> show system configuration rescue
version "7.3"; groups {
    global {
        system {
            host-name router1;
            domain-name customer.net;
            domain-search [ customer.net ];
            backup-router 192.168.124.254;
            name-server {
                172.17.28.11;
                172.17.28.101;
                172.17.28.100;
                172.17.28.10;
            }
            login {
                user regress {
                    uid 928;
                    class ;
                    shell csh;
                    authentication {
                        encrypted-password "$1$kPU..$w.4FGRAGanJ8U4Yq6sbj7."; ##
 SECRET-DATA
                    }
                }
            }
            services {
                ftp;
                rlogin;
                rsh;
                telnet;
            }
        }
    }
....
```

show system rollback

Syntax show system rollback *number*
<compare *number*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display the contents of a previously committed configuration, or the differences between two previously committed configurations.

Options *number*—Number of a configuration to view. The output displays the configuration. The range of values is **0** through **49**.

compare *number* —(Optional) Number of another previously committed (rollback) configuration to compare to rollback ***number***. The output displays the differences between the two configurations. The range of values is **0** through **49**.

Required Privilege Level view

List of Sample Output **show system rollback compare on page 385**

show system rollback compare

```
user@host> show system rollback 3 compare 1
[edit]
+  interfaces {
+      ge-1/1/1 {
+          unit 0 {
+              family inet {
+                  filter {
+                      input mf_plp;
+                  }
+                  address 14.1.1.1/30;
+              }
+          }
+      }
+      ge-1/2/1 {
+          unit 0 {
+              family inet {
+                  filter {
+                      input mf_plp;
+                  }
+                  address 13.1.1.1/30;
+              }
+          }
+      }
+      ge-1/3/0 {
+          unit 0 {
+              family inet {
+                  filter {
+                      input mf_plp;
+                  }
+                  address 12.1.1.1/30;
+              }
+          }
```

```
+         }
+}
```

test configuration

Syntax test configuration *filename*

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Verify that the syntax of a configuration file is correct. If the configuration contains any errors, a message is displayed to indicate the line number and column number in which the error was found.

Options *filename*—Name of the configuration file.

Required Privilege Level view

List of Sample Output **test configuration on page 387**

Output Fields When you enter this command, you are provided feedback on the status of your request.

test configuration

```
user@host> test configuration terminal
[Type ^D to end input]
system {
host-name bluesky;
paris-23;
login;
}
terminal:3:(8) syntax error: paris
  [edit system]
    'paris-23;'
      syntax error
terminal:4:(11) statement must contain additional statements: ;
  [edit system login]
    'login ;'
      statement must contain additional statements
configuration syntax failed
```

PART 8

User and Access Management on J-EX Series Switches

CHAPTER 28

User and Access Management on J-EX Series Switches Overview

J-EX Series Switch Software Features Overview

The following tables list the J-EX Series Switches software features and the Junos OS Release in which they were introduced:

Related Documentation

- J-EX8208 Switch Hardware Overview on page 27
- J-EX8216 Switch Hardware Overview on page 30

Understanding Software Infrastructure and Processes

Each switch runs the Junos OS for J-EX Series Switches on its general-purpose processors. The Junos OS includes processes for Internet Protocol (IP) routing and for managing interfaces, networks, and the chassis.

The Junos OS runs on the Routing Engine. The Routing Engine kernel coordinates communication among the Junos OS processes and provides a link to the Packet Forwarding Engine.

With the J-Web interface and the command-line interface (CLI) to the Junos OS , you configure switching features and routing protocols and set the properties of network interfaces on your switch. After activating a software configuration, use either the J-Web or CLI user interface to monitor the switch, manage operations, and diagnose protocol and network connectivity problems.

- Routing Engine and Packet Forwarding Engine on page 392
- Junos OS Processes on page 392

Routing Engine and Packet Forwarding Engine

A switch has two primary software processing components:

- Packet Forwarding Engine—Processes packets; applies filters, routing policies, and other features; and forwards packets to the next hop along the route to their final destination.
- Routing Engine—Provides three main functions:
  - Creates the packet forwarding switch fabric for the switch, providing route lookup, filtering, and switching on incoming data packets, then directing outbound packets to the appropriate interface for transmission to the network
  - Maintains the routing tables used by the switch and controls the routing protocols that run on the switch.
  - Provides control and monitoring functions for the switch, including controlling power and monitoring system status.

Junos OS Processes

The Junos OS running on the Routing Engine and Packet Forwarding Engine consists of multiple processes that are responsible for individual functions.

The separation of functions provides operational stability, because each process accesses its own protected memory space. In addition, because each process is a separate software package, you can selectively upgrade all or part of the Junos OS, for added flexibility.

Table 55 on page 393 describes the primary Junos OS processes.

Table 55: Junos OS Processes

Process	Name	Description
Chassis process	chassisd	Detects hardware on the system that is used to configure network interfaces. Monitors the physical status of hardware components and field-replaceable units (FRUs), detecting when environment sensors such as temperature sensors are triggered. Relays signals and interrupts—for example, when devices are taken offline, so that the system can close sessions and shut down gracefully.
Ethernet switching process	eswd	Handles Layer 2 switching functionality such as MAC address learning, Spanning Tree protocol and access port security. The process is also responsible for managing Ethernet switching interfaces, VLANs, and VLAN interfaces. Manages Ethernet switching interfaces, VLANs, and VLAN interfaces.
Forwarding process	pfem	Defines how routing protocols operate on the switch. The overall performance of the switch is largely determined by the effectiveness of the forwarding process.
Interface process	dcd	Configures and monitors network interfaces by defining physical characteristics such as link encapsulation, hold times, and keepalive timers.
Management process	mgd	Provides communication between the other processes and an interface to the configuration database. Populates the configuration database with configuration information and retrieves the information when queried by other processes to ensure that the system operates as configured. Interacts with the other processes when commands are issued through one of the user interfaces on the switch. If a process terminates or fails to start when called, the management process attempts to restart it a limited number of times to prevent thrashing and logs any failure information for further investigation.
Routing protocol process	rpd	Defines how routing protocols such as RIP, OSPF, and BGP operate on the device, including selecting routes and maintaining forwarding tables.

Related Documentation

- For more information about processes, see the *Junos OS Network Operations Guide* at http://www.juniper.net/techpubs/software/junos/.
- For more information about basic system parameters, supported protocols, and software processes, see *Junos OS System Basics Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.

CHAPTER 29

User Access Management Configuration

Configuring Management Access for the J-EX Series Switch (J-Web Procedure)

You can manage a J-EX Series switch remotely through the J-Web interface. To communicate with the switch, the J-Web interface uses Hypertext Transfer Protocol (HTTP). HTTP allows easy Web access but no encryption. The data that is transmitted between the Web browser and the switch by means of HTTP is vulnerable to interception and attack. To enable secure Web access the switch supports HTTP over Secure Sockets Layer (HTTPS). You can enable HTTP or HTTPS access on specific interfaces and ports as needed.

Navigate to the Secure Access Configuration page by selecting **Configure** > **System Properties** > **Management Access**. On this page, you can enable HTTP and HTTPS access on interfaces for managing the J-EX Series switch through the J-Web interface. You can also install SSL certificates and enable Junos XML management protocol over SSL with the Secure Access page.

1. Click **Edit** to modify the configuration. Enter information into the Management Access Configuration page as described in Table 56 on page 396.
2. To verify that Web access is enabled correctly, connect to the switch using the appropriate method:
   - For HTTP access—In your Web browser, type **http://*URL*** or **http://*IP address***.
   - For HTTPS access—In your Web browser, type **https://*URL*** or **https://*IP address***.
   - For SSL Junos XML management protocol access—To use this option, you must have a Junos XML management protocol client such as Junos Scope. For information about how to log into Junos Scope, see the *Junos Scope Software User Guide*.

NOTE: After you make changes to the configuration in this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See "Using the Commit Options to Commit Configuration Changes (J-Web Procedure)" on page 334 for details about all commit options.

Table 56: Secure Management Access Configuration Summary

Field	Function	Your Action
Management Access tab		
Management Port IP/Management Port IPv6	Specifies the management port IP address. The software supports both IPv4 (displayed as IP) and IPv6 address. NOTE: IPv6 is not supported on EX 4500 switches.	To specify an IPv4 address: 1. Select the check box **IPv4 address**. 2. Type an IP address—for example: **10.10.10.10**. 3. Enter the subnet mask or address prefix. For example, 24 bits represents **255.255.255.0**. 4. Click **OK**. To specify an IPv6 address: 1. Select the check box **IPv6 address**. 2. Type an IP address—for example:**2001:ab8:85a3::8a2e:370:7334**. 3. Enter the subnet mask or address prefix. 4. Click **OK**.
Default Gateway	Defines a default gateway through which to direct packets addressed to networks that are not explicitly listed in the bridge table constructed by the switch.	For IPv4 address type a 32-bit IP address, in dotted decimal notation. Type a 128-bit IP address for IPv6 address type.
Loopback address	Specifies the IP address of the loopback interface.	Type an IP address.
Subnet Mask	Specifies the subnet mask for the loopback interface.	Enter the subnet mask or address prefix.
Services tab		
Services	Specifies services to be enabled: telnet and SSH.	Select to enable the required services.
Enable Junos XML management protocol over Clear Text	Enables clear text access to the Junos XML management protocol XML scripting API.	To enable clear text access, select the **Enable Junos XML management protocol over Clear Text** check box.
Enable Junos XML protocol over SSL	Enables secure SSL access to the Junos XML management protocol XML scripting API.	To enable SSL access, select the **Enable Junos XML management protocol over SSL** check box.

Table 56: Secure Management Access Configuration Summary *(continued)*

Field	Function	Your Action
Junos XML management protocol Certificate	Specifies SSL certificates to be used for encryption. This field is available only after you create at least one SSL certificate.	To enable an SSL certificate, select a certificate from the Junos XML management protocol SSL Certificate list—for example, **new**.
Enable HTTP	Enables HTTP access on interfaces.	To enable HTTP access, select the **Enable HTTP access** check box. Select and clear interfaces by clicking the direction arrows: • To enable HTTP access on an interface, add the interface to the HTTP Interfaces list. You can either select all interfaces or specific interfaces.
Enable HTTPS	Enables HTTPS access on interfaces.	To enable HTTPS access, select the **Enable HTTPS access** check box. Select and deselect interfaces by clicking the direction arrows: • To enable HTTPS access on an interface, add the interface to the HTTPS Interfaces list. You can either select all interfaces or specific interfaces. NOTE: Specify the certificate to be used for HTTPS access.
Certificates tab		
Certificates	Displays digital certificates required for SSL access to the switch. Allows you to add and delete SSL certificates.	To add a certificate: 1. Have a general SSL certificate available. See Generating SSL Certificates for more information. 2. Click **Add**. The Add a Local Certificate page opens. 3. Type a name in the Certificate Name box—for example, **new**. 4. Open the certificate file and copy its contents. 5. Paste the generated certificate and RSA private key in the Certificate box. To edit a certificate, select it and click **Edit**. To delete a certificate, select it and click **Delete**.

Related Documentation

- Security Features for J-EX Series Switches Overview on page 16

- Understanding J-Web User Interface Sessions on page 133

Generating SSL Certificates to Be Used for Secure Web Access

You can set up secure Web access for a J-EX Series switch. To enable secure Web access, you must generate a digital Secure Sockets Layer (SSL) certificate and then enable HTTPS access on the switch.

To generate an SSL certificate:

1. Enter the following **openssl** command in your SSH command-line interface on a BSD or Linux system on which **openssl** is installed. The **openssl** command generates a self-signed SSL certificate in the privacy-enhanced mail (PEM) format. It writes the certificate and an unencrypted 1024-bit RSA private key to the specified file.

   % openssl req –x509 –nodes –newkey rsa:1024 –keyout *filename*.pem -out *filename*.pem

   where ***filename*** is the name of a file in which you want the SSL certificate to be written—for example, **my-certificate**.

2. When prompted, type the appropriate information in the identification form. For example, type **US** for the country name.

3. Display the contents of the file that you created.

   cat my-certificate.pem

You can use the J-Web Configuration page to install the SSL certificate on the switch. To do this, copy the file containing the certificate from the BSD or Linux system to the switch. Then open the file, copy its contents, and paste them into the Certificate box on the J-Web Secure Access Configuration page.

You can also use the following CLI statement to install the SSL certificate on the switch:

```
[edit]
user@switch# set security certificates local my-signed-cert load-key-file my-certificate.pem
```

Related Documentation

- Configuring Management Access for the J-EX Series Switch (J-Web Procedure) on page 395
- Security Features for J-EX Series Switches Overview on page 16

Configuring MS-CHAPv2 to Provide Password-Change Support (CLI Procedure)

Junos OS for J-EX Series switches enables you to configure the Microsoft Corporation implementation of the Challenge Handshake Authentication Protocol version 2 (MS-CHAPv2) on the switch to provide password-change support. Configuring MS-CHAPv2 on the switch provides users accessing a switch the option of changing the password when the password expires, is reset, or is configured to be changed at next login.

See RFC 2433 at , Microsoft PPP CHAP Extensions, for information about MS-CHAP.

Before you configure MS-CHAPv2 to provide password-change support, ensure that you have:

- Configured RADIUS server authentication. Configure users on the authentication server and set the first-tried option in the authentication order to radius. See “Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch” on page 2267.

To configure MS-CHAPv2, specify the following:

```
[edit system radius-options]
user@switch# set password-protocol mschap-v2
```

You must have the required access permission on the switch in order to change your password.

Related Documentation

- Managing Users (J-Web Procedure) on page 401
- For more about configuring user access, see the *Junos OS Access Privilege Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.

CHAPTER 30

Monitoring Users

Managing Users (J-Web Procedure)

You can use the Users Configuration page for user information to add new users to a J-EX Series switch. For each account, you define a login name and password for the user and specify a login class for access privileges.

To configure users:

1. Select **Configure** > **System Properties** > **User Management**.

   The User Management page displays details of users, the authentication order, the RADIUS servers and TACACS servers present.

2. Click **Edit**.
3. Click any of the following options on the **Users** tab:
   - **Add**—Select this option to add a user. Enter details as described in Table 57 on page 402.
   - **Edit**—Select this option to edit an existing user's details. Enter details as described in Table 57 on page 402.
   - **Delete**—Select this option to delete a user.
4. Click an option on the **Authentication Methods and Order** tab:
   - Authentication Order—Drag and drop the authentication type from the Available Methods section to the Selected Methods. Click the up or down buttons to modify the authentication order.
   - RADIUS server—Click one:
     - **Add**—Select this option to add an authentication server. Enter details as described in Table 58 on page 403.
     - **Edit**—Select this option to modify the authentication server details. Enter details as described in Table 58 on page 403.
     - **Delete**—Select this option to delete an authentication server from the list.

- TACACS server—Click one:
  - **Add**—Select this option to add an authentication server. Enter details as described in Table 58 on page 403.
  - **Edit**—Select this option to modify the authentication server details. Enter details as described in Table 58 on page 403.
  - **Delete**—Select this option to delete an authentication server from the list.

NOTE: After you make changes to the configuration in this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options** > Commit. See "Using the Commit Options to Commit Configuration Changes (J-Web Procedure)" on page 334 for details about all commit options.

Table 57: User Management Configuration Page Summary

Field	Function	Your Action
User Information		
Username (required)	Specifies the name that identifies the user.	Type the username. It must be unique within the switching platform. Do not include spaces, colons, or commas in the username.
User Id	Specifies the user identification.	Type the user's ID.
Full Name	Specifies the user's full name.	Type the user's full name. If the full name contains spaces, enclose it in quotation marks. Do not include colons or commas.
Login Class (required)	Defines the user's access privilege.	Select the user's login class from the list: • **operator** • **read-only** • **super-user/superuser** • **unauthorized** This list also includes any user-defined login classes.
Password	Specifies the login password for this user.	Type the login password for this user. The login password must meet these criteria: • The password must be at least 6 characters long. • It can include alphabetic, numeric, and special characters, but not control characters. • It must contain at least one change of case or character class.
Confirm Password	Verifies the login password for this user.	Retype the login password for this user.

Table 58: Add an Authentication Server

Field	Function	Your Action
IP Address	Specifies the IP address of the server.	Type the server's 32-bit IP address, in dotted decimal notation.
Password	Specifies the password of the server.	Type the password of the server.
Confirm Password	Verifies that the password of the server is entered correctly.	Retype the password of the server.
Server Port	Specifies the port with which the server is associated.	Type the port number.
Source Address	Specifies the source address of the server.	Type the server's 32-bit IP address, in dotted decimal notation.
Retry Attempts	Specifies the number of login retries allowed after a login failure.	Type the number. NOTE: Only 1 retry is permitted for a TACACS server.
Time out	Specifies the time interval to wait before the connection to the server is closed.	Type the interval in seconds.

Related Documentation

- Configuring Management Access for the J-EX Series Switch (J-Web Procedure) on page 395

CHAPTER 31

Troubleshooting User Access Management

Troubleshooting Loss of the Root Password

Problem If you forget the root password for the switch, you can use the password recovery procedure to reset the root password.

NOTE: You need physical access to the switch to recover the root password.

Solution To recover the root password:

1. Power off your switch by unplugging the power cord or turning off the power at the wall switch.
2. Insert one end of the Ethernet cable into the serial port on the management device and connect the other end to the console port on the back of the switch. See Figure 6 on page 405

Figure 6: Connecting to the Console Port on the J-EX Series Switch

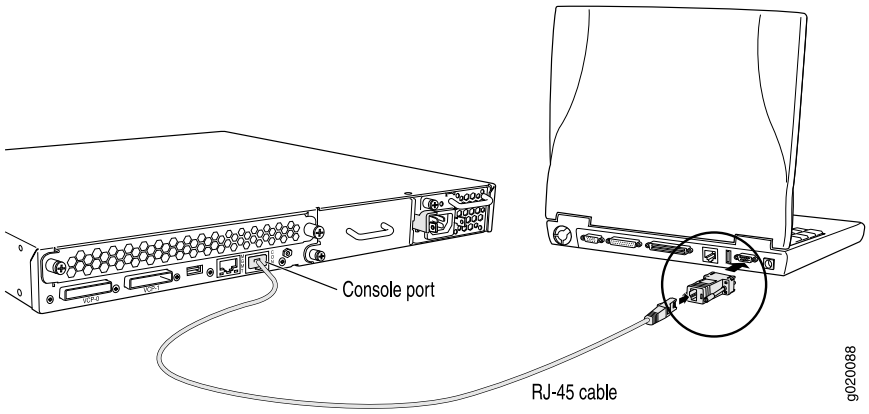

3. On the management device, start your asynchronous terminal emulation application (such as Microsoft Windows Hyperterminal) and select the appropriate **COM** port to use (for example, **COM1**).

4. Configure the port settings as follows:
   - Bits per second: 9600
   - Data bits: 8
   - Parity: None
   - Stop bits: 1
   - Flow control: None
5. Power on your switch by plugging in the power cord or turning on the power at the wall switch.
6. When the following prompt appears, press the Spacebar to access the switch's bootstrap loader command prompt:

   ```
   Hit [Enter] to boot immediately, or space bar for command prompt.
   Booting [kernel] in 1 second...
   ```

7. At the following prompt, type **boot -s** to start up the system in single-user mode:
   loader> boot -s
8. At the following prompt, type **recovery** to start the root password recovery procedure:
   Enter full path name of shell or 'recovery' for root password recovery or RETURN for /bin/sh: recovery

   A series of messages describe consistency checks, mounting of filesystems, and initialization and checkout of management services. Then the CLI prompt appears.
9. Enter configuration mode in the CLI:
   user@switch> configure
10. Set the root password. For example:
    user@switch# set system root-authentication plain-text-password
11. At the following prompt, enter the new root password. For example:
    New password: juniper1

    ```
    Retype new password:
    ```

12. At the second prompt, reenter the new root password.
13. If you are finished configuring the network, commit the configuration.
    root@switch# commit

    ```
    commit complete
    ```

14. Exit configuration mode in the CLI.
    root@switch# exit
15. Exit operational mode in the CLI.
    root@switch> exit
16. At the prompt, enter **y** to reboot the switch.
    Reboot the system? [y/n] y

Related Documentation

- Connecting and Configuring a J-EX Series Switch (CLI Procedure) on page 161
- Connecting and Configuring a J-EX Series Switch (J-Web Procedure) on page 163
- For information about configuring an encrypted root password, configuring SSH keys to authenticate root logins, and configuring special requirements for plain-text passwords, see the *Junos OS System Basics Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.

CHAPTER 32

Configuration Statements for User and Access Management

allow-commands

Syntax allow-commands "*regular-expression*";

Hierarchy Level [edit system login **class** *class-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the operational mode commands that members of a login class can use.

Default If you omit this statement and the **deny-commands** statement, users can issue only those commands for which they have access privileges through the **permissions** statement.

Options ***regular-expression***—Extended (modern) regular expression as defined in POSIX 1003.2. If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Specifying Access Privileges for Junos OS Operational Mode Commands
- **deny-commands on page 415**
- **user on page 432**

allow-configuration

Syntax allow-configuration "*regular-expression*";

Hierarchy Level [edit system login **class** *class-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the configuration mode commands that members of a login class can use.

Default If you omit this statement and the **deny-configuration** statement, users can issue only those commands for which they have access privileges through the **permissions** statement.

Options ***regular-expression***—Extended (modern) regular expression as defined in POSIX 1003.2. If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Specifying Access Privileges for Junos OS Configuration Mode Commands
- Regular Expressions for Allowing and Denying Junos OS Configuration Mode Commands
- **deny-commands on page 415**
- **user on page 432**

announcement

Syntax announcement *text*;

Hierarchy Level [edit system **login**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a system login announcement. This announcement appears after a user logs in.

Options ***text***—Text of the announcement. If the text contains any spaces, enclose it in quotation marks.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration

Related Documentation

- Configuring the Junos OS to Display a System Login Announcement
- **message on page 421**

authentication (Login)

Syntax

```
authentication {
  (encrypted-password "password" | plain-text-password);
  ssh-dsa "public-key";
  ssh-rsa "public-key";
}
```

Hierarchy Level [edit system login user *username*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Authentication methods that a user can use to log in to the router or switch. You can assign multiple authentication methods to a single user.

Options

encrypted-password "*password*"—Message Digest 5 (MD5) or other encrypted authentication. Specify the MD5 or other password. You can specify only one encrypted password for each user.

You cannot configure a blank password for **encrypted-password** using blank quotation marks (" "). You must configure a password whose number of characters range from 1 through 128 characters and enclose the password in quotation marks.

plain-text-password—Plain-text password. The command-line interface (CLI) prompts you for the password and then encrypts it.

ssh-dsa "*public-key*"—SSH version 2 authentication. Specify the SSH public key. You can specify one or more public keys for each user.

ssh-rsa "*public-key*"—SSH version 1 and SSH version 2 authentication. Specify the SSH public key. You can specify one or more public keys for each user.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Junos OS User Accounts
- **root-authentication on page 426**

authentication-order

Syntax authentication-order [*authentication-methods*];

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the order in which the software tries different user authentication methods when attempting to authenticate a user. For each login attempt, the software tries the authentication methods in order, starting with the first one, until the password matches.

Default If you do not include the **authentication-order** statement, users are verified based on their configured passwords.

Options ***authentication-methods***—One or more authentication methods, listed in the order in which they should be tried. The method can be one or more of the following:

- **password**—Use the password configured for the user with the **authentication** statement at the **[edit system login user]** hierarchy level.
- **radius**—Use RADIUS authentication services.
- **tacplus**—Use TACACS+ authentication services.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring the Junos OS Authentication Order for RADIUS, TACACS+, and Local Password Authentication
- **authentication on page 411**

change-type

Syntax change-type (character-sets | set-transitions);

Hierarchy Level [edit system login **password**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set requirements for using character sets in plain-text passwords. When you combine this statement with the **minimum-changes** statement, you can check for the total number of character sets included in the password or for the total number of character-set changes in the password. Newly created passwords must meet these requirements.

Options Specify one of the following:

- **character-sets**—The number of character sets in the password. Valid character sets include uppercase letters, lowercase letters, numbers, punctuation, and other special characters.
- **set-transitions**—The number of transitions between character sets.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Special Requirements for Junos OS Plain-Text Passwords
- **minimum-changes on page 422**

class (Assigning a Class to an Individual User)

Syntax class *class-name*;

Hierarchy Level [edit system login **user** *username*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a user's login class. You must configure one class for each user.

Options ***class-name***—One of the classes defined at the **[edit system login class]** hierarchy level.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Junos OS User Accounts

class (Defining Login Classes)

Syntax

```
class class-name {
  allow-commands "regular-expression";
  allow-configuration "regular-expression";
  deny-commands "regular-expression";
  deny-configuration "regular-expression";
  idle-timeout minutes;
  permissions [ permissions ];
}
```

Hierarchy Level [edit system login]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define a login class.

Options *class-name*—A name you choose for the login class.

The remaining statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Defining Junos OS Login Classes
- user on page 432

deny-commands

Syntax deny-commands "*regular-expression*";

Hierarchy Level [edit system login **class**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the operational mode commands that the user is denied permission to issue even though the permissions set with the **permissions** statement would allow it.

Default If you omit this statement and the **allow-commands** statement, users can issue only those commands for which they have access privileges through the **permissions** statement.

Options ***regular-expression***—Extended (modern) regular expression as defined in POSIX 1003.2. If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Specifying Access Privileges for Junos OS Operational Mode Commands
- **allow-commands on page 409**
- **user on page 432**

deny-configuration

Syntax deny-configuration "*regular-expression*";

Hierarchy Level [edit system login **class**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the configuration mode commands that the user is denied permission to issue even though the permissions set with the **permissions** statement would allow it.

Default If you omit this statement and the **allow-configuration** statement, users can issue only those commands for which they have access privileges through the **permissions** statement.

Options ***regular-expression***—Extended (modern) regular expression as defined in POSIX 1003.2. If the regular expression contains any spaces, operators, or wildcard characters, enclose it in quotation marks.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Specifying Access Privileges for Junos OS Operational Mode Commands
- **allow-configuration on page 410**
- **user on page 432**

format

Syntax format (des | md5 | sha1);

Hierarchy Level [edit system login password]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the authentication algorithm for plain-text passwords.

Default For the Junos OS, the default encryption format is **md5**. For Junos-FIPS software, the default encryption format is **sha1**.

Options The hash algorithm that authenticates the password can be one of three algorithms:

- **des**—Has a block size of 8 bytes; its key size is 48 bits long.
- **md5**—Produces a 128-bit digest.
- **sha1**—Produces a 160-bit digest.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Special Requirements for Junos OS Plain-Text Passwords

full-name

Syntax full-name *complete-name*;

Hierarchy Level [edit system login **user**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the complete name of a user.

Options ***complete-name***—Full name of the user. If the name contains spaces, enclose it in quotation marks.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring Junos OS User Accounts

idle-timeout

Syntax idle-timeout *minutes*;

Hierarchy Level [edit system login **class** *class-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For a login class, configure the maximum time that a session can be idle before the user is logged off the router or switch. The session times out after remaining at the CLI operational mode prompt for the specified time.

Default If you omit this statement, a user is never forced off the system after extended idle times.

Options ***minutes***—Maximum idle time.

Range: 0 through 100,000 minutes

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring the Timeout Value for Idle Login Sessions
- **user on page 432**

login

Syntax

```
login {
  announcement text;
  class class-name {
    allow-commands "regular-expression";
    allow-configuration "regular-expression";
    deny-commands "regular-expression";
    deny-configuration "regular-expression";
    idle-timeout minutes;
    login-tip;
    permissions [ permissions ];
  }
  message text;
  password {
    change-type (set-transitions | character-set);
    format (md5 | sha1 | des);
    maximum-length length;
    minimum-changes number;
    minimum-length length;
  }
  retry-options {
    backoff-threshold number;
    backoff-factor seconds;
    minimum-time seconds;
    tries-before-disconnect number;
  }
  user username {
    full-name complete-name;
    uid uid-value;
    class class-name;
    authentication authentication;
    (encrypted-password "password" | plain-text-password);
    ssh-rsa "public-key";
    ssh-dsa "public-key";
  }
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure user access to the router or switch.

Options The remaining statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Defining Junos OS Login Classes

login-alarms

Syntax login-alarms;

Hierarchy Level [edit system login class admin]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For J-EX Series switches only. Show system alarms automatically when an **admin** user logs on to the router.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

login-tip

Syntax login-tip;

Hierarchy Level [edit system login class *class-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable CLI tips at login.

Default Disabled.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring CLI Tips

maximum-length

Syntax maximum-length *length*;

Hierarchy Level [edit system login passwords]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the maximum number of characters allowed in plain-text passwords. Newly created passwords must meet this requirement.

Default For Junos-FIPS software, the maximum number of characters for plain-text passwords is **20**. For Junos OS, no maximum is set.

Options **length**—The maximum number of characters the password can include.
Range: 1 to 64 characters

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Special Requirements for Junos OS Plain-Text Passwords

message

Syntax message *text*;

Hierarchy Level [edit system login]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a system login message. This message appears before a user logs in.

Options ***text***—Text of the message.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration

Related Documentation
- Configuring the Junos OS to Display a System Login Message
- **announcement on page 410**

minimum-changes

Syntax minimum-changes *number*;

Hierarchy Level [edit system login passwords]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the minimum number of character sets (or character set changes) required in plain-text passwords. Newly created passwords must meet this requirement.

This statement is used in combination with the **change-type** statement. If the change-type is **character-sets**, then the number of character sets included in the password is checked against the specified minimum. If change-type is **set-transitions**, then the number of character set changes in the password is checked against the specified minimum.

Default For Junos OS, the minimum number of changes is 1. For Junos-FIPS Software, the minimum number of changes is 3.

Options *number*—The minimum number of character sets (or character set changes) required for the password.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Special Requirements for Junos OS Plain-Text Passwords
- **change-type on page 413**

minimum-length

Syntax minimum-length *length*;

Hierarchy Level [edit system login passwords]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the minimum number of characters required in plain-text passwords. Newly created passwords must meet this requirement.

Default For Junos OS, the minimum number of characters for plain-text passwords is six. For Junos-FIPS software, the minimum number of characters for plain-text passwords is 10.

Options **length**—The minimum number of characters the password must include.

Range: 6 to 20 characters

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Special Requirements for Junos OS Plain-Text Passwords
- **maximum-length on page 421**

password (Login)

Syntax
```
password {
  change-type (set-transitions | character-set);
  format (md5 | sha1 | des);
  maximum-length length;
  minimum-changes number;
  minimum-length length;
}
```

Hierarchy Level [edit system login]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure special requirements such as character length and encryption format for plain-text passwords. Newly created passwords must meet these requirements.

The remaining statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Special Requirements for Junos OS Plain-Text Passwords
- **maximum-length on page 421**

permissions

Syntax permissions [*permissions*];

Hierarchy Level [edit system login **class**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the login access privileges to be provided on the router or switch.

Options ***permissions***—Privilege type. For a list of permission flag types, see Junos OS Access Privilege Levels Overview.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring Access Privilege Levels
- **user on page 432**

radius-options

Syntax
```
radius-options {
  attributes {
    nas-ip-address ip-address;
  }
  password-protocol mschap-v2;
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure RADIUS options for the NAS-IP address for outgoing RADIUS packets and password protocol used in RADIUS packets.

Options ***ip-address***—IP address of the network access server (NAS) that requests user authentication.

mschap-v2—Protocol MS-CHAPv2, used for password authentication and password changing.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring RADIUS Authentication

retry-options

Syntax

```
retry-options {
  backoff-threshold number;
  backoff-factor seconds;
  maximum-time seconds;
  minimum-time seconds;
  tries-before-disconnect number;
}
```

Hierarchy Level [edit system login]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Maximum number of times a user can attempt to enter a password while logging in through SSH or Telnet before being disconnected.

Options

backoff-threshold *number*—Threshold for the number of failed login attempts before the user experiences a delay when attempting to reenter a password. Use the **backoff-factor** option to specify the length of delay, in seconds.

Range: 1 through 3

Default: 2

backoff-factor *seconds*—Length of delay after each failed login attempt. The length of delay increases by this value for each subsequent login attempt after the value specified in the **backoff-threshold** option.

Range: 5 through 10

Default: 5

maximum-time *seconds*—Maximum length of time that the connection remains open for the user to enter a username and password to log in. If the user remains idle and does not enter a username and password within the configured **maximum-time**, the connection is closed.

Range: 20 through 300

Default: 120

minimum-time *seconds*—Minimum length of time that the connection remains open while the user is attempting to enter a password to log in.

Range: 20 through 60

Default: 20

tries-before-disconnect *number*—Maximum number of times a user is allowed to attempt to enter a password to log in through SSH or Telnet.

Range: 1 through 10

Default: 10

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Limiting the Number of User Login Attempts for SSH and Telnet Sessions
- **rate-limit on page 487**

root-authentication

Syntax

```
root-authentication {
  (encrypted-password "password" | plain-text-password);
  ssh-dsa "public-key";
  ssh-rsa "public-key";
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the authentication methods for the root-level user, whose username is **root**.

Options

encrypted-password "*password*"— MD5 or other encrypted authentication. Specify the MD5 or other password. You can specify only one encrypted password.

You cannot configure a blank password for **encrypted-password** using blank quotation marks (" "). You must configure a password whose number of characters range from 1 through 128 characters and enclose the password in quotation marks.

plain-text-password—Plain-text password. The CLI prompts you for the password and then encrypts it. The CLI displays the encrypted version, and the software places the encrypted version in its user database. You can specify only one plain-text password.

ssh-dsa "*public-key*"—SSH version 2 authentication. Specify the DSA (SSH version 2) public key. You can specify one or more public keys.

ssh-rsa "*public-key*"—SSH version 1 authentication. Specify the RSA (SSH version 1 and SSH version 2) public key. You can specify one or more public keys.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring the Root Password
- **authentication on page 411**

root-login

Syntax root-login (allow | deny | deny-password);

Hierarchy Level [edit system services ssh]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Control user access through SSH.

Default Allow user access through SSH.

Options

allow—Allow users to log in to the router or switch as root through SSH.

deny—Disable users from logging in to the router or switch as root through SSH.

deny-password—Allow users to log in to the router or switch as root through SSH when the authentication method (for example, RSA authentication) does not require a password.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring SSH Service for Remote Access to the Router or Switch
- Configuring SSH Service for Remote Access to the Router or Switch

tacplus-options

Syntax

```
tacplus-options {
  service-name service-name;
  (no-cmd-attribute-value | exclude-cmd-attribute);
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure TACACS+ options for authentication and accounting.

Options

service-name *service-name*—The name of the authentication service used when configuring multiple TACACS+ servers to use the same authentication service.

Default: junos-exec

no-cmd-attribute-value—Set the cmd attribute value to an empty string in the TACACS+ accounting start and stop requests to enable logging of accounting records in the correct log file on a TACACS+ server.

exclude-cmd-attribute—Exclude the cmd attribute value completely from start and stop accounting records to enable logging of accounting records in the correct log file on a TACACS+ server.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring TACACS+ Authentication
- Configuring TACACS+ System Accounting
- Junos OS Authentication Order for RADIUS, TACACS+, and Password Authentication

tacplus-server

Syntax

```
tacplus-server server-address {
    secret password;
    single-connection;
    source-address source-address;
    timeout seconds;
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the TACACS+ server.

Options ***server-address***—Address of the TACACS+ authentication server.

The remaining statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring TACACS+ Authentication

traceoptions (Address-Assignment Pool)

Syntax

```
traceoptions {
  file filename {
    files number;
    size maximum-file-size;
    match regex;
    <world-readable | no-world-readable>;
  }
  flag address-assignment;
  flag all;
  flag configuration;
  flag framework;
  flag ldap;
  flag local-authentication;
  flag radius;
}
```

Hierarchy Level [edit system processes general-authentication-service]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure tracing options.

Options

file *filename*—Name of the file that receives the output of the tracing operation. Enclose the name in quotation marks. All files are placed in the directory **/var/log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

Range: 2 through 1000

Default: 3 files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **address-assignment**—All address-assignment events
- **all**—All tracing operations
- **configuration**—Configuration events
- **framework**—Authentication framework events
- **ldap**—LDAP authentication events
- **local-authentication**—Local authentication events
- **radius**—RADIUS authentication events

match *regex*—(Optional) Refine the output to include lines that contain the regular expression.

no-world-readable—(Optional) Restrict access to the originator of the trace operation only.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and filename.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: **10** KB through **1** GB

Default: **128** KB

world-readable—(Optional) Enable unrestricted file access.

Required Privilege Level
admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Tracing Address-Assignment Pool Processes
- Configuring Address-Assignment Pools

uid

Syntax
uid *uid-value*;

Hierarchy Level
[edit system login **user**]

Release Information
Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Configure a user identifier for a login account.

Options
uid-value—Number associated with the login account. This value must be unique on the router or switch.

Range: **100** through **64000**

Required Privilege Level
admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring Junos OS User Accounts

user (Access)

Syntax

```
user username {
  authentication {
  class class-name;
  (encrypted-password "password" | plain-text-password);
  full-name complete-name;
  ssh-dsa "public-key";
  ssh-rsa "public-key";
  uid uid-value;
  }
}
```

Hierarchy Level [edit system login]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure access permission for individual users.

Options The remaining statements are explained separately.

Required Privilege Level
admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Junos OS User Accounts
- **class on page 413**

CHAPTER 33

Operational Mode Commands for User and Access Management

request message

Syntax
request message all message "*text*"
request message message "*text*" (terminal *terminal-name* | user *user-name*)

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display a message on the screens of all users who are logged in to the router or switch or on specific screens.

Options
all—Display a message on the terminal of all users who are currently logged in.

message "*text*"—Message to display.

terminal *terminal-name*—Name of the terminal on which to display the message.

user *user-name*—Name of the user to whom to direct the message.

Required Privilege Level
maintenance

List of Sample Output
request message message on page 434

Output Fields
When you enter this command, you are provided feedback on the status of your request.

request message message

```
user@host> request message message "Maintenance window in 10 minutes" user maria
Message from user@host on ttyp0 at 20:27 ...
Maintenance window in 10 minutes
EOF
```

show subscribers

Syntax

show subscribers
<address *address*>
<client-type *client-type*>
<interface *interface*>
<logical-system *logical-system*>
<mac-address *mac-address*>
<profile-name *profile-name*>
<routing-instance *routing-instance*>
<stacked-vlan-id *stacked-vlan-id*>
<subscriber-state *subscriber-state*>
<vlan-id *vlan-id*>
<count | detail | extensive |summary (all | logical-system *logical-system* | routing-instance *routing-instance*) | terse>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.
client-type, **mac-address**, **subscriber-state**, **extensive**, and **summary** options introduced in Junos OS Release 10.2.
count option usage with other options introduced in Junos OS Release 10.2

Description

Display information for active subscribers.

Options

address—(Optional) Display subscribers whose IP address matches the specified address.

client-type—(Optional) Display subscribers whose client type matches the specified client type (DHCP, L2TP, PPP, PPPOE, or VLAN).

count—(Optional) Display the count of total subscribers and active subscribers for any specified option. You can use the count option alone or with the **address**, **client-type**, **interface**, **logical-system**, **mac-address**, **profile-name**, **routing-instance**, **stacked-vlan-id**, **subscriber-state**, and **vlan-id** options.

interface—(Optional) Display subscribers whose interface matches the specified interface.

logical system—(Optional) Display subscribers whose logical system matches the specified logical system.

mac-address—(Optional) Display subscribers whose MAC address matches the specified MAC address.

profile name—(Optional) Display subscribers whose dynamic profile matches the specified profile name.

routing instance—(Optional) Display subscribers whose routing instance matches the specified routing instance.

subscriber-state—(Optional) Display subscribers whose subscriber state matches the specified subscriber state (ACTIVE, CONFIGURED, INIT, TERMINATED, or TERMINATING).

vlan-id—(Optional) Display subscribers whose VLAN ID matches the specified VLAN ID.

stacked-vlan-id—(Optional) Display subscribers whose stacked VLAN ID matches the specified stacked VLAN ID.

detail | terse | extensive—(Optional) Display the specified level of output.

summary—(Optional) Display summary output.

NOTE: Due to display limitations, logical system and routing instance output values are truncated when necessary.

Required Privilege Level view

List of Sample Output

show subscribers on page 438
show subscribers detail (IPv4) on page 438
show subscribers detail (IPv6) on page 438
show subscribers logical-system on page 439
show subscribers count on page 439
show subscribers routing-instance inst1 count on page 439
show subscribers vlan-id on page 439
show subscribers vlan-id detail on page 439
show subscribers stacked-vlan-id detail on page 439
show subscribers stacked-vlan-id vlan-id detail (Combined Output) on page 439
show subscribers stacked-vlan-id vlan-id interface detail (Combined Output for a Specific Interface) on page 440
show subscribers client-type dhcp detail on page 440
show subscribers extensive on page 440
show subscribers summary on page 440
show subscribers summary all on page 441
show subscribers terse on page 441

Output Fields Table 59 on page 436 lists the output fields for the **show subscribers** command. Output fields are listed in the approximate order in which they appear.

Table 59: show subscribers Output Fields

Field Name	Field Description
User Name	Name of subscriber.
Type	Subscriber client type (DHCP, VLAN, PPP, PPPOE, or L2TP).
IP Address	Subscriber IPv4 address.
IP Netmask	Subscriber IP netmask.
IPv6 Address	Subscriber IPv6 address.
IPv6 Prefix	Subscriber IPv6 prefix.

Table 59: show subscribers Output Fields *(continued)*

Field Name	Field Description
IPv6 Prefix Length	Length of the subscriber IPv6 prefix.
Logical System	Logical system associated with the subscriber.
Routing Instance	Routing instance associated with the subscriber.
Interface	Interface associated with the subscriber. The router displays subscribers whose interface matches or begins with the specified interface.
Interface Type	Whether the subscriber interface is static or dynamic.
Dynamic Profile Name	Dynamic profile used for the subscriber.
MAC Address	MAC address associated with the subscriber.
State	Current state of the subscriber session (Init, Configured, Active, Terminating, Terminated).
VLAN Id	VLAN ID associated with the subscriber in the form *tpid.vlan-id*.
Stacked VLAN Id	Stacked VLAN ID associated with the subscriber in the form *tpid.vlan-id*.
RADIUS Accounting ID	RADIUS accounting ID associated with the subscriber.
Agent Circuit ID	Option 82 agent circuit ID associated with the subscriber.
Agent Remote ID	Option 82 agent remote ID associated with the subscriber.
DHCP Relay IP Address	IP address used by the DHCP relay agent.
Login Time	Date and time at which the subscriber logged in.
Session ID	ID number for a subscriber service session.
Service Sessions	Number of service sessions (that is, a service activated using RADIUS CoA) associated with the subscribers.
Service Session Name	Service session profile name.
IPv4 Input Filter Name	Name assigned to the IPv4 input filter (client or service session).
IPv4 Output Filter Name	Name assigned to the IPv4 output filter (client or service session).
IPv6 Input Filter Name	Name assigned to the IPv6 input filter (client or service session).
IPv6 Output Filter Name	Name assigned to the IPv6 output filter (client or service session).
IFL Input Filter Name	Name assigned to the logical interface input filter (client or service session).

Table 59: show subscribers Output Fields *(continued)*

Field Name	Field Description
IFL Output Filter Name	Name assigned to the logical interface output filter (client or service session).
Subscribers by State	Number of subscribers summarized by state. The summary information includes the following: • Init—Number of subscriber currently in the initialization state. • Configured—Number of configured subscribers. • Active—Number of active subscribers. • Terminating—Number of subscribers currently terminating. • Terminated—Number of terminated subscribers. Summary information includes subscriber counts per state and the total number of subscribers.
Subscribers by Client Type	Number of subscribers summarized by client type. Client types can include DHCP, VLAN, PPP, PPPOE, and L2TP. Summary information includes subscriber counts per client type and the total number of subscribers.
Subscribers by LS:RI	Number of subscribers summarized by logical system:routing instance (LS:RI) combination. Summary information includes subscriber counts per LS:RI and the total number of subscribers.

show subscribers

```
user@host> show subscribers
Interface           IP Address/VLAN ID   User Name         LS:RI
ge-1/3/0.1073741824   100                                   default:default
demux0.1073741824     100.0.0.10           WHOLESALER-CLIENT default:default
demux0.1073741825     101.0.0.3            RETAILER1-CLIENT  test1:retailer1
demux0.1073741826     102.0.0.3            RETAILER2-CLIENT  test1:retailer2
```

show subscribers detail (IPv4)

```
user@host> show subscribers detail
Type: DHCP
IP Address: 100.20.9.7
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073744127
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:10:95:00:00:98
State: Active
Radius Accounting ID: jnpr :2304
Login Time: 2009-08-25 14:43:52 PDT
Service Sessions: 2
```

show subscribers detail (IPv6)

```
user@host> show subscribers detail
Type: DHCP
IPv6 Address: 1080:0:0:0:8:800:200C:417A
Ipv6 Prefix: fec0:1:1:1::/128
Logical System: default1
Routing Instance: default
Interface: demux0.1073744127
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:10:95:00:00:98
State: Active
```

```
Radius Accounting ID: jnpr :2304
Login Time: 2009-08-25 14:43:52 PDT
Service Sessions: 2
```

show subscribers logical-system

```
user@host> show subscribers logical-system test1 terse
Interface           IP Address/VLAN ID   User Name         LS:RI
demux0.1073741825     101.0.0.3            RETAILER1-CLIENT  test1:retailer1
demux0.1073741826     102.0.0.3            RETAILER2-CLIENT  test1:retailer2
```

show subscribers count

```
user@host> show subscribers count
Total Subscribers: 188, Active Subscribers: 188
```

show subscribers routing-instance inst1 count

```
user@host> show subscribers routing-instance inst1 count
Total Subscribers: 188, Active Subscribers: 183
```

show subscribers vlan-id

```
user@host> show subscribers vlan-id 100
Interface           IP Address                                User Name
ge-1/0/0.1073741824
ge-1/2/0.1073741825
```

show subscribers vlan-id detail

```
user@host> show subscribers vlan-id 100 detail
Type: VLAN
Interface: ge-1/0/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: vlan-prof-tpid
State: Active
VLAN Id: 100
Login Time: 2009-03-11 06:48:54 PDT

Type: VLAN
Interface: ge-1/2/0.1073741825
Interface type: Dynamic
Dynamic Profile Name: vlan-prof-tpid
State: Active
VLAN Id: 100
Login Time: 2009-03-11 06:48:54 PDT
```

show subscribers stacked-vlan-id detail

```
user@host> show subscribers stacked-vlan-id 101 detail
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT
```

show subscribers stacked-vlan-id vlan-id detail (Combined Output)

```
user@host> show subscribers stacked-vlan-id 101 vlan-id 100 detail
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT
```

show subscribers stacked-vlan-id vlan-id interface detail (Combined Output for a Specific Interface)

```
user@host> show subscribers stacked-vlan-id 101 vlan-id 100 interface ge-1/2/0.* detail
Type: VLAN
Interface: ge-1/2/0.1073741824
Interface type: Dynamic
Dynamic Profile Name: svlan-prof
State: Active
Stacked VLAN Id: 0x8100.101
VLAN Id: 0x8100.100
Login Time: 2009-03-27 11:57:19 PDT
```

show subscribers client-type dhcp detail

```
user@host> show subscribers client-type dhcp detail
Type: DHCP
IP Address: 100.20.9.7
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073744127
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:10:95:00:00:98
State: Active
Radius Accounting ID: jnpr :2304
Login Time: 2009-08-25 14:43:52 PDT

Type: DHCP
IP Address: 100.20.10.7
IP Netmask: 255.255.0.0
Logical System: default
Routing Instance: default
Interface: demux0.1073744383
Interface type: Dynamic
Dynamic Profile Name: dhcp-demux-prof
MAC Address: 00:10:94:00:01:f3
State: Active
Radius Accounting ID: jnpr :2560
Login Time: 2009-08-25 14:43:56 PDT
```

show subscribers extensive

```
user@host> show subscribers extensive
Type: DHCP
       IPv6 Prefix: 2001::40:0:0:0/74
       IPv6 Prefix Length: 64
       Logical System: default
           Routing Instance: default
           Interface: demux0.1073741825
           Interface type: Dynamic
           Dynamic Profile Name: dhcp-demux-prof
           State: Active
           Radius Accounting ID: jnpr :2
           Agent Circuit ID: abc
           Remote Circuit ID: xyz
           Login Time: 2010-03-31 14:27:19 PDT
           Service Sessions: 1
           IPv6 Input Filter Name: demux0-inet6-in
              Session ID: 213
              Service Session Name: service-profile
              IPv6 Input Filter Name: dfwd1-demux.1073741825-in
```

show subscribers summary

```
user@host> show subscribers summary

Subscribers by State
```

```
 Init      3
 Configured    2
 Active   183
 Terminating    2
 Terminated    1

 TOTAL        191

Subscribers by Client Type
 DHCP     107
 PPP       76
 VLAN       8

 TOTAL        191
```

show subscribers summary all

```
user@host> show subscribers summary all
Subscribers by State
 Init         3
 Configured       2
 Active      183
 Terminating       2
 Terminated       1

 TOTAL          191

Subscribers by Client Type
 DHCP      107
 PPP          76
 VLAN          8

 TOTAL         191

Subscribers by LS:RI
 default:default   1
 default:ri1     28
 default:ri2     16
 ls1:default    22
 ls1:riA     38
 ls1:riB     44
 logsysX:routinstY  42

 TOTAL     191
```

show subscribers terse

```
user@host> show subscribers summary terse
Interface           IP Address/VLAN ID   User Name         LS:RI
ge-1/3/0.1073741824   100                                   default:default
demux0.1073741824     100.0.0.10           WHOLESALER-CLIENT default:default
demux0.1073741825     101.0.0.3            RETAILER1-CLIENT  test1:retailer1
demux0.1073741826     102.0.0.3            RETAILER2-CLIENT  test1:retailer2
```

PART 9

Junos OS for J-EX Series Switches System Services

CHAPTER 34

System Services Overview

DHCP Overview

DHCP Services for J-EX Series Switches Overview

A Dynamic Host Configuration Protocol (DHCP) server can automatically allocate IP addresses and also deliver configuration settings to client hosts on a subnet.

DHCP is particularly useful for managing a pool of IP addresses among hosts. An IP address can be leased to a host for a limited period of time, allowing the DHCP server to share a limited number of IP addresses among a group of hosts that do not need permanent IP addresses.

DHCP, through the use of the automatic software download feature, can also be used to install software packages on J-EX Series Switches. Users can define a path to a software package on the DHCP server, and then the DHCP server communicates this path to J-EX Series switches acting as DHCP clients as part of the DHCP message exchange process. The DHCP clients that have been configured for automatic software download receive these messages and, when the software package name in the DHCP server message is different from that of the software package that booted the DHCP client switch, download and install the software package. See “Upgrading Software Using Automatic Software Download on J-EX Series Switches” on page 82.

To configure DHCP access service for a J-EX Series switch, you can use either the Junos OS command-line interface (CLI) or the J-Web user interface.

For detailed information about configuring DHCP services, see the *Junos OS System Basics Configuration Guide* at http://www.juniper.net/techpubs/software/junos/. The configuration for DHCP service on the J-EX Series switch includes the **dhcp** statement at the **[edit system services]** hierarchy level.

You can monitor DHCP services for the switch by using either operational-mode CLI commands or the J-Web interface.

Related Documentation

- For information about configuring DHCP services with the CLI, see the *Junos OS System Basics Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.
- Configuring DHCP Services (J-Web Procedure) on page 447
- Upgrading Software Using Automatic Software Download on J-EX Series Switches on page 82
- Monitoring DHCP Services on page 451

DHCP/BOOTP Relay for J-EX Series Switches Overview

You can configure the J-EX Series Switch to act as a Dynamic Host Configuration Protocol (DHCP) or Bootstrap Protocol (BOOTP) relay agent. This means that a locally attached host can issue a DHCP or BOOTP request as a broadcast message. If the switch sees this broadcast message, it relays the message to a specified DHCP or BOOTP server. You should configure the switch to be a DHCP/BOOTP relay agent if you have locally attached hosts and a distant DHCP or BOOTP server.

For detailed information about configuring a DHCP/BOOTP relay agent, see the *Junos OS Policy Framework Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.

You can configure a J-EX Series Switch to use the gateway IP address (giaddr) as the source IP address of the switch for relayed DHCP packets when the switch is used as the DHCP relay agent. For information on configuring this option, see the **source-address-giaddr** configuration statement.

NOTE: Because DHCP/BOOTP messages are broadcast and are not directed to a specific server, switch, or router, J-EX Series switches cannot function as both a DHCP server and a DHCP/BOOTP relay agent at the same time. The Junos OS generates a commit error if both options are configured at the same time, and the commit will not succeed until one of the options is removed.

Related Documentation

- For information about configuring the switch as a DHCP/BOOTP relay agent, see the *Junos OS Policy Framework Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.
- DHCP Services for J-EX Series Switches Overview on page 445

CHAPTER 35

System Services Configuration

Configuring DHCP Services (J-Web Procedure)

Use the J-Web DHCP Configuration pages to configure DHCP pools for subnets and static bindings for DHCP clients on a J-EX Series switch. If DHCP pools or static bindings are already configured, use the Configure Global DHCP Parameters Configuration page to add settings for these pools and static bindings. Settings that have been previously configured for DHCP pools or static bindings are not overridden when you use the Configure Global DHCP Parameters Configuration page.

To configure the DHCP server:

1. Select **Configure** > **Services** > **DHCP**.
2. Access a DHCP Configuration page:
   - To configure a DHCP pool for a subnet, click **Add** in the DHCP Pools box.
   - To configure a static binding for a DHCP client, click **Add** in the DHCP Static Binding box.
   - To globally configure settings for existing DHCP pools and static bindings, click **Configure Global DHCP Parameters**.
3. Enter information into the DHCP Configuration pages as described in Table 60 on page 448.
4. To apply the configuration, click **Apply**.

NOTE: After you make changes to the configuration in this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select Commit Options > Commit. See "Using the Commit Options to Commit Configuration Changes (J-Web Procedure)" on page 334 for details about all commit options.

Table 60: DHCP Server Configuration Pages Summary

Field	Function	Your Action
DHCP Pool Information		
DHCP Subnet (required)	Specifies the subnet on which DHCP is configured.	Type an IP address prefix.
Address Range (Low) (required)	Specifies the lowest address in the IP address pool range.	Type an IP address that is part of the subnet specified in DHCP Subnet.
Address Range (High) (required)	Specifies the highest address in the IP address pool range.	Type an IP address that is part of the subnet specified in DHCP Subnet. This address must be greater than the address specified in Address Range (Low).
Exclude Addresses	Specifies addresses to exclude from the IP address pool.	• To add an excluded address, type the address next to the **Add** button, and click **Add**. • To delete an excluded address, select the address in the Exclude Addresses box, and click **Delete**.
Lease Time		
Maximum Lease Time (Seconds)	Specifies the maximum length of time a client can hold a lease. (Dynamic BOOTP lease lengths can exceed this maximum time.)	Type a number from 60 through 4,294,967,295 (seconds). You can also type **infinite** to specify a lease that never expires.
Default Lease Time (Seconds)	Specifies the length of time a client can hold a lease for clients that do not request a specific lease length.	Type a number from 60 through 2,147,483,647 (seconds). You can also type **infinite** to specify a lease that never expires.
Server Information		
Server Identifier	Specifies the IP address of the DHCP server reported to a client.	Type the IP address of the server. If you do not specify a server identifier, the primary address of the interface on which the DHCP exchange occurs is used.
Domain Name	Specifies the domain name that clients must use to resolve hostnames.	Type the name of the domain.
Domain Search	Specifies the order—from top to bottom—in which clients must append domain names when resolving hostnames using DNS.	• To add a domain name, type the name next to the **Add** button, and click **Add**. • To delete a domain name, select the name in the Domain Search box, and click **Delete**.
DNS Name Servers	Defines a list of DNS servers the client can use, in the specified order—from top to bottom.	• To add a DNS server, type an IP address next to the **Add** button, and click **Add**. • To remove a DNS server, select the IP address in the DNS Name Servers box, and click **Delete**.

Table 60: DHCP Server Configuration Pages Summary *(continued)*

Field	Function	Your Action
Gateway Routers	Defines a list of relay agents on the subnet, in the specified order—from top to bottom.	• To add a relay agent, type an IP address next to the **Add** button, and click **Add**. • To remove a relay agent, select the IP address in the Gateway Routers box, and click **Delete**.
WINS Servers	Defines a list of NetBIOS name servers, in the specified order—from top to bottom.	• To add a NetBIOS name server, type an IP address next to the **Add** button, and click **Add**. • To remove a NetBIOS name server, select the IP address in the WINS Servers box, and click **Delete**.
Boot Options		
Boot File	Specifies the path and filename of the initial boot file to be used by the client.	Type a path and filename.
Boot Server	Specifies the TFTP server that provides the initial boot file to the client.	Type the IP address or hostname of the TFTP server.
DHCP Static Binding Information		
DHCP MAC Address (required)	Specifies the MAC address of the client to be permanently assigned a static IP address.	Type the hexadecimal MAC address of the client.
Fixed IP Addresses (required)	Defines a list of IP addresses permanently assigned to the client. A static binding must have at least one fixed address assigned to it, but multiple addresses are also allowed.	• To add an IP address, type it next to the **Add** button, and click **Add**. • To remove an IP address, select it in the Fixed IP Addresses box, and click **Delete**.
Host Name	Specifies the name of the client used in DHCP messages exchanged between the server and the client. The name must be unique to the client within the subnet on which the client resides.	Type a client hostname.
Client Identifier	Specifies the name of the client used by the DHCP server to index its database of address bindings. The name must be unique to the client within the subnet on which the client resides.	Type a client identifier in string form.
Hexadecimal Client Identifier	Specifies the name of the client, in hexadecimal form, used by the DHCP server to index its database of address bindings. The name must be unique to the client within the subnet on which the client resides.	Type a client identifier in hexadecimal form.

Related Documentation

- DHCP Services for J-EX Series Switches Overview on page 445
- Monitoring DHCP Services on page 451

Configuring a DHCP SIP Server (CLI Procedure)

You can use the **sip-server** statement on the J-EX Series switch to configure option 120 on a DHCP server. The DHCP server sends configured option values—Session Initiation Protocol (SIP) server addresses or names—to DHCP clients when they request them. Previously, you were only allowed to specify a SIP server by address using **[edit system services dhcp option 120]**. You specify either an IPv4 address or a fully qualified domain name to be used by SIP clients to locate a SIP server. You cannot specify both an address and name in the same statement.

To configure a SIP server using the ***address*** option:

```
[edit system services dhcp]
user@switch# set sip-server address
```

For example, to configure one address:

```
[edit system services dhcp]
user@switch set sip-server 172.168.0.11
```

To configure a SIP server using the ***name*** option:

```
[edit system services dhcp]
user@switch# set sip-server name
```

For example, to configure a name:

```
[edit system services dhcp]
user@switch set sip-server abc.example.com
```

Related Documentation

- DHCP Services for J-EX Series Switches Overview on page 445
- *Junos OS System Basics Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

CHAPTER 36

Monitoring System Services

Monitoring DHCP Services

Purpose A switch or router can operate as a DHCP server. Use the monitoring functionality to view information about dynamic and static DHCP leases, conflicts, pools, and statistics.

Action To monitor the DHCP server in the J-Web interface, select **Monitor > Services > DHCP**.

To monitor the DHCP server in the CLI, enter the following CLI commands:

- **show system services dhcp binding**
- **show system services dhcp conflict**
- **show system services dhcp pool**
- **show system services dhcp statistics**
- **show system services dhcp relay-statistics**
- **show system services dhcp global**
- **show system services dhcp client**

Meaning Table 61 on page 452 summarizes the output fields in DHCP displays in the J-Web interface.

Table 61: Summary of DHCP Output Fields

Field	Values	Additional Information
Global tab		
Name	This column displays the following information: • Boot lease length • Domain Name • Name servers • Server identifier • Domain search • Gateway routers • WINS server • Boot file • Boot server • Default lease time • Minimum lease time • Maximum lease time	
Value	Displays the value for each of the parameters in the Name column.	
Bindings tab		
Allocated Address	List of IP addresses the DHCP server has assigned to clients.	
MAC Address	Corresponding media access control (MAC) address of the client.	
Binding Type	Type of binding assigned to the client: **dynamic** or **static**.	DHCP servers can assign a dynamic binding from a pool of IP addresses or a static binding to one or more specific IP addresses.
Lease Expires	Date and time the lease expires, or **never** for leases that do not expire.	
Pools tab		
Pool Name	Subnet on which the IP address pool is defined.	
Low Address	Lowest address in the IP address pool.	
High Address	Highest address in the IP address pool.	
Excluded Addresses	Addresses excluded from the address pool.	
Clients tab		

Table 61: Summary of DHCP Output Fields *(continued)*

Field	Values	Additional Information
Interface Name	Name of the logical interface.	
Hardware Address	Vendor identification.	
Status	State of the client binding.	
Address Obtained	IP address obtained from the DHCP server.	
Update Server	Indicates whether server update is enabled.	
Lease Obtained	Date and time the lease was obtained.	
Lease Expires	Date and time the lease expires.	
Renew	Reacquires an IP address from the server for the interface. When you click this option, the command sends a discover message if the client state is INIT and a renew request message if the client state is BOUND. For all other states it performs no action.	
Release	Clears other resources received earlier from the server, and reinitializes the client state to INIT for the particular interface.	
Conflicts tab		
Detection Time	Date and time the client detected the conflict.	
Detection Method	How the conflict was detected.	Only client-detected conflicts are displayed.
Address	IP address where the conflict occurs.	The addresses in the conflicts list remain excluded until you use the **clear system services dhcp conflict** command to manually clear the list.
DHCP Statistics		
Relay Statistics tab		
Packet Counters	Displays the number of packet counters.	
Dropped Packet Counters	Graphically displays the number of dropped packet counters.	

Table 61: Summary of DHCP Output Fields *(continued)*

Field	Values	Additional Information
Statistics tab		
Packets dropped	Total number of packets dropped and the number of packets dropped due to a particular condition.	
Messages received	Number of BOOTREQUEST, DHCPDECLINE, DHCPDISCOVER, DHCPINFORM, DHCPRELEASE, and DHCPREQUEST messages sent from DHCP clients and received by the DHCP server.	
Messages sent	Number of BOOTREPLY, DHCPACK, DHCPOFFER, and DHCPNAK messages sent from the DHCP server to DHCP clients.	

Related Documentation

- Configuring DHCP Services (J-Web Procedure) on page 447
- DHCP Services for J-EX Series Switches Overview on page 445

CHAPTER 37

Configuration Statements for System Services

boot-file

Syntax boot-file *filename*;

Hierarchy Level [edit system services **dhcp**],
[edit system services dhcp **pool**],
[edit system services dhcp **static-binding**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For J-EX Series switches only. Set the boot file advertised to DHCP clients. After the client receives an IP address and the boot file location from the DHCP server, the client uses the boot image stored in the boot file to complete DHCP setup.

Options ***filename***—The location of the boot file on the boot server. The filename can include a pathname.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring the Router, Switch, or Interface to Act as a DHCP Server on J Series Services Routers and J-EX Series Switches
- **boot-server on page 456**

boot-server (DHCP)

Syntax boot-server *address*;

Hierarchy Level [edit system services **dhcp**],
[edit system services dhcp **pool**],
[edit system services dhcp **static-binding**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For J-EX Series switches only. Configure the name of the boot server advertised to DHCP clients. The client uses a boot file located on the boot server to complete DHCP setup.

Options ***address***—Address of a boot server. You must specify an IPv4 address, not a hostname.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- **boot-file on page 455**

bootp

Syntax

```
bootp {
  client-response-ttl number;
  description text-description;
  interface interface-group {
    client-response-ttl number;
    description text-description;
    maximum-hop-count number;
    minimum-wait-time seconds;
    no-listen;
    server address {
      <logical-system logical-system-name> <routing-instance
        [ routing-instance-names ]>;
    }
  }
  maximum-hop-count number;
  minimum-wait-time seconds;
  server address {
    <logical-system logical-system-name> <routing-instance
      [ routing-instance-names ]>;
  }
}
```

Hierarchy Level [edit forwarding-options helpers]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configures a router, switch, or interface to act as a Dynamic Host Configuration Protocol (DHCP) or bootstrap protocol (BOOTP) relay agent.

DHCP relaying is disabled.

Options The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents
- Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 2646

ca-name

Syntax ca-name *ca-identity*;

Hierarchy Level [edit security certificates **certification-authority**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description (Encryption interface on J-EX Series switches) Specify the certificate authority (CA) identity to use in the certificate request.

Options ***ca-identity***—CA identity to use in the certificate request.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring Digital Certificates for an ES PIC

cache-size

Syntax cache-size *bytes*;

Hierarchy Level [edit security **certificates**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description (Encryption interface on J-EX Series switches) Configure the cache size for digital certificates.

Options ***bytes***—Cache size for digital certificates.

Range: 64 through 4,294,967,295

Default: 2 megabytes (MB)

NOTE: We recommend that you limit your cache size to 4 MB.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration

Related Documentation
- Configuring Digital Certificates for an ES PIC

cache-timeout-negative

Syntax cache-timeout-negative *seconds*;

Hierarchy Level [edit security **certificates**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description (Encryption interface on J-EX Series switches) Configure a negative cache for digital certificates.

Options ***seconds***—Negative time to cache digital certificates, in seconds.
Range: 10 through 4,294,967,295
Default: 20

CAUTION: Configuring a large negative cache value can lead to a denial-of-service attack.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration

certificates

Syntax

```
certificates {
  cache-size bytes;
  cache-timeout-negative seconds;
  certification-authority ca-profile-name {
    ca-name ca-identity;
    crl file-name;
    encoding (binary | pem);
    enrollment-url url-name;
    file certificate-filename;
    ldap-url url-name;
  }
  enrollment-retry attempts;
  local certificate-name {
    certificate-key-string;
    load-key-file URL-or-path;
  }
  maximum-certificates number;
  path-length certificate-path-length;
}
```

Hierarchy Level [edit security]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description (Encryption interface J-EX Series switches) Configure the digital certificates for IPsec.

The remaining statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

certification-authority

Syntax

```
certification-authority ca-profile-name {
    ca-name ca-identity;
    crl file-name;
    encoding (binary | pem);
    enrollment-url url-name;
    file certificate-filename;
    ldap-url url-name;
}
```

Hierarchy Level [edit security **certificates**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description (Encryption interface on J-EX Series switches) Configure a certificate authority profile name.

The remaining statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration

client-identifier

Syntax client-identifier (ascii *client-id* | hexadecimal *client-id*);

Hierarchy Level [edit system services dhcp **static-binding**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For J-EX Series switches only. Configure the client's unique identifier. This identifier is used by the DHCP server to index its database of address bindings. Either a client identifier or the client's MAC address is required to uniquely identify the client on the network.

Options ***client-id***—A name or number that uniquely identifies the client on the network. The client identifier can be an ASCII string or hexadecimal digits.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

connection-limit

Syntax connection-limit *limit*;

Hierarchy Level [edit system services finger],
[edit system services **ftp**],
[edit system services **ssh**],
[edit system services **telnet**],
[edit system services xnm-clear-text],
[edit system services xnm-ssl]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the maximum number of established connections for each type of system service (finger, ftp, ssh, telnet, xnm-clear-text, or xnm-ssl) for each IP protocol, such as IPv6 and IPv4.

Options *limit*—(Optional) Maximum number of established connections.

Range: 1 through 250

Default: 75

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring clear-text or SSL Service for Junos XML Management Protocol Client Applications
- Configuring DTCP-over-SSH Service for the Flow-Tap Application
- Configuring Finger Service for Remote Access to the Router
- Configuring FTP Service for Remote Access to the Router or Switch
- Configuring SSH Service for Remote Access to the Router or Switch
- Configuring Telnet Service for Remote Access to a Router

crl (Encryption Interface)

Syntax crl *file-name*;

Hierarchy Level [edit security **certificates**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description (Encryption interface on J-EX Series switches) Configure the certificate revocation list (CRL). A CRL is a time-stamped list identifying revoked certificates, which is signed by a CA and made available to the participating IPsec peers on a regular periodic basis.

Options ***file-name***—Specify the file from which to read the CRL.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration

default-lease-time

Syntax default-lease-time *seconds*;

Hierarchy Level [edit system services **dhcp**],
[edit system services dhcp **pool**],
[edit system services dhcp **static-binding**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For J-EX Series switches only. Specify the length of time in seconds that a client holds the lease for an IP address assigned by a DHCP server. This setting is used if a lease time is not requested by the client.

Options ***seconds***—Number of seconds the lease can be held.

Default: **86400** (1day)

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- **maximum-lease-time on page 479**

description

Syntax description *text-description*;

Hierarchy Level [edit forwarding-options **helpers bootp**],
[edit forwarding-options **helpers bootpinterface** *interface-group*],
[edit forwarding-options **helpers domain**],
[edit forwarding-options **helpers domain interface** *interface-name*],
[edit forwarding-options **helpers tftp**],
[edit forwarding-options **helpers tftpinterface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Describe a BOOTP, DHCP, Domain Name System (DNS), or Trivial File Transfer Protocol (TFTP) service, or an interface that is configured for the service.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring DNS and TFTP Packet Forwarding
- Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents

dhcp

Syntax

```
dhcp {
  boot-file filename;
  boot-server (address | hostname);
  default-lease-time seconds;
  domain-name domain-name;
  domain-search [domain-list];
  maximum-lease-time seconds;
  name-server {
    address;
  }
  option {
    [ (id-number option-type option-value) | (id-number array option-type option-value) ];
  }
  pool address/prefix-length {
    address-range {
      low address;
      high address;
    }
    exclude-address {
      address;
    }
  }
  router {
    address;
  }
  static-binding mac-address {
    fixed-address {
      address;
    }
    host hostname;
    client-identifier (ascii client-id | hexadecimal client-id);
  }
  server-identifier address;
  wins-server {
    address;
  }
}
```

Hierarchy Level [edit system services]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For J-EX Series switches only. Configure a router, switch, or interface as a DHCP server. A DHCP server can allocate network addresses and deliver configuration information to client hosts on a TCP/IP network.

The remaining statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- System Management Configuration Statements

domain

Syntax

```
domain {
  description text-description;
  interface interface-name {
    broadcast;
    description text-description;
    no-listen;
    server address <logical-system logical-system-name> <routing-instance
      routing-instance-name>;
  }
  server address <logical-system logical-system-name> <routing-instance
    routing-instance-name>;
}
```

Hierarchy Level [edit forwarding-options **helpers**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable DNS request packet forwarding.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring DNS and TFTP Packet Forwarding

domain-name (DHCP)

Syntax domain-name *domain-name*;

Hierarchy Level [edit system services **dhcp**],
[edit system services dhcp **pool**],
[edit system services dhcp **static-binding**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For J-EX Series switches only. Configure the name of the domain in which clients search for a DHCP server host. This is the default domain name that is appended to hostnames that are not fully qualified.

Options *domain-name*—Name of the domain.

domain-search

Syntax domain-search [*domain-list*];

Hierarchy Level [edit system],
[edit system services dhcp],
[edit system services dhcp pool],
[edit system services dhcp static-binding]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a list of domains to be searched.

Options ***domain-list***—A list of domain names to search. The list can contain up to six domain names, with a total of up to 256 characters.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring the Domains to Search When a Router or Switch Is Included in Multiple Domains

encoding

Syntax encoding (binary | pem);

Hierarchy Level [edit security ike policy *ike-peer-address*],
[edit security certificates **certification-authority** *ca-profile-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description (Encryption interface on J-EX Series switches) Specify the file format used for the **local-certificate** and **local-key-pair** statements.

Options **binary**—Binary file format.

pem—Privacy-enhanced mail (PEM), an ASCII base 64 encoded format.

Default: binary

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

enrollment-retry

Syntax enrollment-retry *attempts*;

Hierarchy Level [edit security **certificates**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description (Encryption interface on J-EX Series switches) Specify how many times a router or switch can resend a digital certificate request.

Options ***attempts***—Number of enrollment retries.
Range: **0** through **100**
Default: **0**

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

enrollment-url

Syntax enrollment-url *url-name*;

Hierarchy Level [edit security certificates **certification-authority** *ca-profile-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description (Encryption interface on J-EX Series switches) Specify where your router or switch sends Simple Certificate Enrollment Protocol-based (SCEP-based) certificate enrollment requests (certificate authority URL).

Options ***url-name***—Certificate authority URL.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

file

Syntax file *certificate-filename*;

Hierarchy Level [edit security certificates **certification-authority** *ca-profile-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description (Encryption interface on J-EX Series switches) Specify the file from which to read the digital certificate.

Options ***certificate-filename***—File from which to read the digital certificate.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

ftp

Syntax
```
ftp {
    connection-limit limit;
    rate-limit limit;
}
```

Hierarchy Level [edit system **services**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Allow FTP requests from remote systems to the local router or switch.

Options The remaining statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

helpers

Syntax

```
helpers {
  bootp {
    client-response-ttl number;
    description text-description;
    interface interface-group {
      client-response-ttl number;
      description text-description;
      maximum-hop-count number;
      minimum-wait-time seconds;
      no-listen;
      server address {
        <logical-system logical-system-name>
          <routing-instance [ routing-instance-names ]>;
      }
    }
    maximum-hop-count number;
    minimum-wait-time seconds;
    server address {
      <logical-system logical-system-name>
        <routing-instance [ routing-instance-names ]>;
    }
  }
  domain {
    description text-description;
    interface interface-name {
      broadcast;
      description text-description;
      no-listen;
      server address <logical-system logical-system-name> <routing-instance
        routing-instance-name>;
    }
    server address <logical-system logical-system-name> <routing-instance
      routing-instance-name>;
  }
  port port-number {
    description text-description;
    interface interface-name {
      broadcast;
      description text-description;
      no-listen;
      server address <logical-system logical-system-name> <routing-instance
        routing-instance-name>;
    }
    server address <logical-system logical-system-name> <routing-instance
      routing-instance-name>;
  }
  tftp {
    description text-description;
    interface interface-name {
      broadcast;
      description text-description;
      no-listen;
```

```
        server address <logical-system logical-system-name> <routing-instance
          routing-instance-name>;
      }
      server address <logical-system logical-system-name> <routing-instance
        routing-instance-name>;
    }
    traceoptions {
      file filename <files number> <match regular-expression> <size bytes> <world-readable |
        no-world-readable>;
      flag flag;
      level level;
      no-remote-trace level;
    }
}
```

Hierarchy Level [edit forwarding-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable TFTP or DNS request packet forwarding, or configure the router, switch, or interface to act as a DHCP/BOOTP relay agent. Use only one server address per interface or global configuration.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring DNS and TFTP Packet Forwarding

http

Syntax

```
http {
  interfaces [ interface-names ];
  port port;
}
```

Hierarchy Level [edit system **services web-management**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the port and interfaces for HTTP service, which is unencrypted.

Options

interfaces [*interface-names*]—Name of one or more interfaces on which to allow the HTTP service. By default, HTTP access is allowed through built-in Fast Ethernet or Gigabit Ethernet interfaces only.

The remaining statement is explained separately.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Management Access for the J-EX Series Switch (J-Web Procedure) on page 395
- *J-Web Interface User Guide*
- **https on page 473**
- **port on page 486**
- **web-management on page 506**

https

Syntax

```
https {
  interfaces [ interface-names ];
  local-certificate name;
  port port;
}
```

Hierarchy Level [edit system **services web-management**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the secure version of HTTP (HTTPS) service, which is encrypted.

Options

interfaces [***interface-names*****]**—Name of one or more interfaces on which to allow the HTTPS service. By default, HTTPS access is allowed through any ingress interface, but HTTP access is allowed through built-in Fast Ethernet or Gigabit Ethernet interfaces only.

local-certificate ***name***—Name of the X.509 certificate for a Secure Sockets Layer (SSL) connection. An SSL connection is configured at the **[edit security certificates local]** hierarchy.

The remaining statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Management Access for the J-EX Series Switch (J-Web Procedure) on page 395
- *J-Web Interface User Guide*
- **http on page 472**
- **port on page 486**
- **web-management on page 506**

interface (BOOTP)

Syntax

```
interface interface-group {
  client-response-ttl number;
  description text-description;
  maximum-hop-count number;
  minimum-wait-time seconds;
  no-listen;
  server address {
    <logical-system logical-system-name> <routing-instance [ routing-instance-names ]>;
  }
}
```

Hierarchy Level [edit forwarding-options **helpers bootp**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the interface for a DHCP and BOOTP relay agent.

Options ***interface-group***—Sets a logical interface or group of logical interfaces with a specific DHCP relay configuration.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents
- Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 2646

interface (DNS and TFTP Packet Forwarding or Relay Agent)

Syntax

```
interface interface-name {
  broadcast;
  description text-description;
  no-listen;
  server address <logical-system logical-system-name> <routing-instance
    routing-instance-name>;
}
```

Hierarchy Level [edit forwarding-options **helpers domain**],
[edit forwarding-options **helpers tftp**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the interface for monitoring and forwarding DNS or TFTP requests.

Options ***interface-name***—Name of the interface.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring DNS and TFTP Packet Forwarding

ldap-url

Syntax <ldap-url *url-name*>;

Hierarchy Level [edit security certificates **certification-authority** *ca-profile-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description (Encryption interface on J-EX Series switches) (Optional) Specify the Lightweight Directory Access Protocol (LDAP) URL for digital certificates.

Options ***url-name***—Name of the LDAP URL.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

load-key-file

Syntax load-key-file;

Hierarchy Level [edit system root-authentication],
[edit system login user *username* authentication]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Load RSA (SSH version 1 and SSH version 2) and DSA (SSH version 2) public keys from a file. The file is a URL containing one or more SSH keys.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring the Root Password
- Configuring Junos OS User Accounts

local

Syntax

```
local certificate-name {
    certificate-key-string;
    load-key-file URL-or-path;
}
```

Hierarchy Level [edit security certificates]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Import a paired X.509 private key and authentication certificate, to enable Junos XML management protocol client applications to establish Secure Sockets Layer (SSL) connections to the router or switch.

Options

certificate-key-string—String of alphanumeric characters that constitute the private key and certificate.

certificate-name—Name that uniquely identifies the certificate.

load-key-file *URL-or-path*—File that contains the private key and certificate. It can be one of two types of values:

- Pathname of a file on the local disk (assuming you have already used another method to copy the certificate file to the router's or switch's local disk)
- URL to the certificate file location (for instance, on the computer where the Junos XML management protocol client application runs)

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Importing SSL Certificates for Junos XML Management Protocol Support

local-certificate

Syntax local-certificate;

Hierarchy Level [edit system services service-deployment],
[edit system services web-management https],
[edit system services xnm-ssl]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Import or reference an SSL certificate.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring clear-text or SSL Service for Junos XML Management Protocol Client Applications
- Generating SSL Certificates to Be Used for Secure Web Access on page 398
- Importing SSL Certificates for Junos XML Management Protocol Support

maximum-certificates

Syntax maximum-certificates *number*;

Hierarchy Level [edit security certificates]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description (Encryption interface on J-EX Series switches) Configure the maximum number of peer digital certificates to be cached.

Options *number*—Maximum number of peer digital certificates to be cached.
Range: 64 through 4,294,967,295 peer certificates
Default: 1024 peer certificates

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

maximum-hop-count

Syntax maximum-hop-count *number*;

Hierarchy Level [edit forwarding-options **helpers bootp**],
[edit forwarding-options **helpers bootpinterface** *interface-group*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the maximum number of hops allowed.

Options ***number***—Maximum number of hops.
Default: 4 hops

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

maximum-lease-time

Syntax maximum-lease-time *seconds*;

Hierarchy Level [edit system services dhcp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For J-EX Series switches. Specify the maximum length of time in seconds for which a client can request and hold a lease on a DHCP server.

An exception is that the dynamic BOOTP lease length can exceed the maximum lease length specified.

Options ***seconds***—The maximum number of seconds the lease can be held.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration

Related Documentation

- **default-lease-time on page 463**

minimum-wait-time

Syntax minimum-wait-time *seconds*;

Hierarchy Level [edit forwarding-options **helpers bootp**],
[edit forwarding-options **helpers bootpinterface** *interface-group*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the minimum time allowed.

Options ***seconds***—Minimum time.

Default: 0 seconds

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

name-server

Syntax
```
name-server {
  address;
}
```

Hierarchy Level [edit system],
[edit system services dhcp],
[edit system services dhcp pool],
[edit system services dhcp static-binding]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure one or more Domain Name System (DNS) name servers.

Options ***address***—Address of the name server. To configure multiple name servers, include multiple ***address*** options.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring a DNS Name Server for Resolving a Hostname into Addresses

no-listen

Syntax no-listen;

Hierarchy Level [edit forwarding-options **helpers bootp interface***interface-group*],
[edit forwarding-options **helpers domain interface** *interface-name*],
[edit forwarding-options **helpers tftp interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable recognition of DNS requests or stop packets from being forwarded on a logical interface, a group of logical interfaces, a router, or a switch.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring DNS and TFTP Packet Forwarding

outbound-ssh

Syntax

```
[edit system services]
outbound-ssh {
  client client-id {
    address {
      port port-number;
      retry number;
      timeout seconds;
    }
    device-id device-id;
    keep-alive {
      retry number;
      timeout seconds;
    }
    reconnect-strategy (in-order | sticky);
    secret password;
    services netconf;
  }
  traceoptions {
    file filename <files number> <match regex> <size size> <world-readable |
      no-world-readable>;
    flag flag;
    no-remote-trace;
  }
}
```

Hierarchy Level [edit system services]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a router or switch running the Junos OS behind a firewall to communicate with client management applications on the other side of the firewall.

Default To configure transmission of the router's or switch's device ID to the application, include the **device-id** statement at the **[edit system services]** hierarchy level.

Options

client-id—Identifies the **outbound-ssh** configuration stanza on the router or switch. Each **outbound-ssh** stanza represents a single outbound SSH connection. This attribute is not sent to the client.

device-id—Identifies the router or switch to the client during the initiation sequence.

keep-alive—(Optional) When configured, specifies that the router or switch send keepalive messages to the management server. To configure the keepalive message, you must set both the **timeout** and **retry** attributes.

reconnect-strategy—(Optional) Specify the method the router or switch uses to reestablish a disconnected outbound SSH connection. Two methods are available:

- **in-order**—Specify that the router or switch first attempt to establish an outbound SSH session based on the management server address list. The router or switch attempts

to establish a session with the first server on the list. If this connection is not available, the router or switch attempts to establish a session with the next server, and so on down the list until a connection is established.

- **sticky**—Specify that the router or switch first attempt to reconnect to the management server that it was last connected to. If the connection is unavailable, it attempts to establish a connection with the next client on the list and so forth until a connection is made.

retry—Number of keepalive messages the router or switch sends without receiving a response from the client before the current SSH connection is disconnected. The default is three messages.

secret—(Optional) Router's or switch's public SSH host key. If added to the **outbound-ssh** statement, during the initialization of the outbound SSH service, the router or switch passes its public key to the management server. This is the recommended method of maintaining a current copy of the router's or switch's public key.

timeout—Length of time that the Junos OS server waits for data before sending a keep alive signal. The default is 15 seconds.

When reconnecting to a client, the router or switch attempts to reconnect to the client based on the **retry** and **timeout** values for each client listed.

address—Hostname or the IPv4 address of the NSM application server. You can list multiple clients by adding each client's IP address or hostname along with the following connection parameters:

- **port**—Outbound SSH port for the client. The default is port 22.
- **retry**—Number of times the router or switch attempts to establish an outbound SSH connection before giving up. The default is three tries.
- **timeout**—Length of time that the router or switch attempts to establish an outbound SSH connection before giving up. The default is fifteen seconds.

filename—(Optional) By default, the filename of the log file used to record the trace options is the name of the traced process (for example, **mib2d** or **snmpd**). Use this option to override the default value.

files—(Optional) Maximum number of trace files generated. By default, the maximum number of trace files is 10. Use this option to override the default value.

When a trace file reaches its maximum size, the system archives the file and starts a new file. The system archives trace files by appending a number to the filename in sequential order from 1 to the maximum value (specified by the default value or the options value set here). Once the maximum value is reached, the numbering sequence is restarted at 1, overwriting the older file.

size—(Optional) Maximum size of the trace file in kilobytes (KB). Once the maximum file size is reached, the system archives the file. The default value is 1000 KB. Use this option to override the default value.

match—(Optional) When used, the system only adds lines to the trace file that match the regular expression specified. For example, if the match value is set to **=error**, the system only records lines to the trace file that include the string **error**.

services—Services available for the session. Currently, NETCONF is the only service available.

world-readable | no-world-readable—(Optional) Whether the files are accessible by the originator of the trace operation only or by any user. By default, log files are only accessible by the user that started the trace operation (**no-world-readable**).

all | configuration | connectivity—(Optional) Type of tracing operation to perform.

all—Log all events.

configuration—Log all events pertaining to the configuration of the router or switch.

connectivity—Log all events pertaining to the establishment of a connection between the client server and the router or switch.

no-remote-trace—(Optional) Disable remote tracing.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Outbound SSH Service
- System Management Configuration Statements

path-length

Syntax
path-length *certificate-path-length*;

Hierarchy Level
[edit security **certificates**]

Release Information
Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
(Encryption interface on J-EX Series switches) Configure the digital certificate path length.

Options
certificate-path-length—Digital certificate path length.
Range: 2 through 15 certificates
Default: 15 certificates

Required Privilege Level
admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

pool

Syntax

```
pool address/prefix-length {
  address-range {
    low address;
    high address;
  }
  exclude-address {
    address;
  }
}
```

Hierarchy Level [edit system services dhcp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For J-EX Series switches . Configure a pool of IP addresses for DHCP clients on a subnet. When a client joins the network, the DHCP server dynamically allocates an IP address from this pool.

Options

address-range—Lowest and highest IP addresses in the pool that are available for dynamic address assignment. If no range is specified, the pool will use all available addresses within the subnet specified. (Broadcast addresses, interface addresses, and excluded addresses are not available.)

exclude-address—Addresses within the range that are not used for dynamic address assignment. You can exclude one or more addresses within the range.

The remaining statements are explained separately.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

port (HTTP/HTTPS)

Syntax port *port-number*;

Hierarchy Level [edit system **services web-management**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the port on which the HTTP or HTTPS service is connected.

Options ***port-number***—The TCP port number on which the specified service listens.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Table 56 on page 396
- *J-Web Interface User Guide*
- **http on page 472**
- **https on page 473**
- **web-management on page 506**

port (SRC Server)

Syntax port *port-number*;

Hierarchy Level [edit system services service-deployment servers *server-address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the port number on which to contact the SRC server.

Options ***port-number***—(Optional) The TCP port number for the SRC server.
Default: 3333

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring the Junos OS to Work with SRC Software

protocol-version

Syntax protocol-version *version*;

Hierarchy Level [edit system services ssh]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the secure shell (SSH) protocol version.

Options ***version***—SSH protocol version
Values: v1, u2, or [v1 v2]
Default: [v1 v2]

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring SSH Service for Remote Access to the Router or Switch

rate-limit

Syntax rate-limit *limit*;

Hierarchy Level [edit system services finger],
[edit system services ftp],
[edit system services ssh],
[edit system services telnet],
[edit system services xnm-clear-text],
[edit system services xnm-ssl]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Maximum number of connection attempts on an access service.

Options **rate-limit *limit***—(Optional) Maximum number of connection attempts allowed per minute.
Range: 1 through 250
Default: 150

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring clear-text or SSL Service for Junos XML Management Protocol Client Applications

server (DHCP and BOOTP Relay Agent)

Syntax

server *address* {
  <logical-system *logical-system-name*> <routing-instance [*routing-instance-names*]>;
}

Hierarchy Level

[edit forwarding-options **helpers bootp**],
[edit forwarding-options **helpers bootp interface** *interface-group*]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure the router or switch to act as a DHCP and BOOTP relay agent.

Options

- ***address***—One or more addresses of the server.
- **logical-system *logical-system-name***—(Optional) Logical system of the server.
- **routing-instance [*routing-instance-names*]**—(Optional) Routing instance name or names that belong to the DHCP or BOOTP relay agent.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Routers, Switches, and Interfaces as DHCP and BOOTP Relay Agents

server (DNS and TFTP Service)

Syntax server *address* <logical-system *logical-system-name*> <routing-instance *routing-instance-name*>;

Hierarchy Level [edit forwarding-options **helpers domain**],
[edit forwarding-options **helpers domain interface** *interface-name*],
[edit forwarding-options **helpers tftp**],
[edit forwarding-options **helpers tftp interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the DNS or TFTP server for forwarding DNS or TFTP requests. Only one server can be specified for each interface.

Options ***address***—Address of the server.

logical-system *logical-system-name*—(Optional) Logical system of the server.

routing-instance [*routing-instance-names*]—(Optional) Set the routing instance name or names that belong to the DNS server or TFTP server.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring DNS and TFTP Packet Forwarding

server-identifier

Syntax server-identifier *address*;

Hierarchy Level [edit system services dhcp],
[edit system services dhcp pool],
[edit system services dhcp static-binding]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For J-EX Series switches . Configure a server identifier. The identifier can be used to identify a DHCP server in a DHCP message. It can also be used as a destination address from clients to servers (for example, when the boot file is set, but not the boot server).

Servers include the server identifier in **DHCPOFFER** messages so that clients can distinguish between multiple lease offers. Clients include the server identifier in **DHCPREQUEST** messages to select a lease and indicate which offer is accepted from multiple lease offers. Also, clients can use the server identifier to send unicast request messages to specific DHCP servers to renew a current lease.

This address must be a manually assigned, static IP address. The server cannot send a request and receive an IP address from itself or another DHCP server.

Default If no server identifier is set, the DHCP server sets the server identifier based on the primary interface address used by the server to receive a client request. For example, if the client sends a DHCP request and the server receives it on **fe-0/0/0** and the primary interface address is **1.1.1.1**, then the server identifier is set to **1.1.1.1**.

Options ***address***—IPv4 address of the server. This address must be accessible by all clients served within a specified range of addresses (based on an address pool or static binding).

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

servers

Syntax

```
servers server-address {
    port port-number;
}
```

Hierarchy Level [edit system services service-deployment]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure an IPv4 address for the Session and Resource Control (SRC) server.

Options *server-address*—The TCP port number.

Default: 3333

The remaining statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring the Junos OS to Work with SRC Software

service-deployment

Syntax

```
service-deployment {
    servers server-address {
        port port-number;
    }
    source-address source-address;
}
```

Hierarchy Level [edit system services]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable Junos OS to work with the Session and Resource Control (SRC) software.

The remaining statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring the Junos OS to Work with SRC Software

services

Syntax

```
services {
  dhcp {
    dhcp_services;
  }
  finger {
    <connection-limit limit>;
    <rate-limit limit>;
  }
  ftp {
    <connection-limit limit>;
    <rate-limit limit>;
  }
  ssh {
    protocol-version [v1 v2];
    <connection-limit limit>;
    <rate-limit limit >;
    root-login (allow | deny | deny-password);
  }
  service-deployment {
    servers server-address {
      port-number port-number;
    }
    source-address source-address;
  }
  telnet {
    <connection-limit limit>;
    <rate-limit limit>;
  }
  web-management {
    http {
      interfaces [ interface-names ];
      port port;
    }
    https {
      interfaces [ interface-names ];
      local-certificate name;
      port port;
    }
    session {
      idle-timeout [ minutes ];
      session-limit [ session-limit ];
    }
  }
  xnm-clear-text {
    <connection-limit limit>;
    <rate-limit limit>;
  }
  xnm-ssl {
    <connection-limit limit>;
    <local-certificate name>
    <rate-limit limit>;
  }
```

}

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the router or switch so that users on remote systems can access the local router or switch through the DHCP server, finger, rlogin, SSH, telnet, Web management, Junos XML management protocol clear-text, Junos XML management protocol SSL, and network utilities or enable Junos OS to work with the Session and Resource Control (SRC) software.

The remaining statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring clear-text or SSL Service for Junos XML Management Protocol Client Applications
- Configuring Junos OS to Work with SRC Software

session

Syntax

```
session {
  idle-timeout [ minutes ];
  session-limit [ session-limit ];
}
```

Hierarchy Level [edit system **services web-management**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure limits for the number of minutes a session can be idle before it times out, and configure the number of simultaneous J-Web user login sessions.

Options

idle-timeout *minutes*—Configure the number of minutes a session can be idle before it times out.

Range: 1 through 1440

Default: 1440

session-limit *session-limit*—Configure the maximum number of simultaneous J-Web user login sessions.

Range: 1 through 1024

Default: Unlimited

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- *J-Web Interface User Guide*

sip-server

Syntax sip-server [*address* | *name*];

Hierarchy Level [edit system **services dhcp**],
[edit system **services dhcp pool**],
[edit system **services dhcp static-binding**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure Session Initiation Protocol (SIP) server addresses or names for DHCP servers.

Options ***address***—IPv4 address of the SIP server. To configure multiple SIP servers, include multiple ***address*** options. This address must be accessible by all clients served within a specified range of addresses (based on an address pool or static binding).

name—Fully qualified domain name of the SIP server. To configure multiple SIP servers, include multiple ***name*** options. This domain name must be accessible by all clients served within a specified range of addresses (based on an address pool or static binding).

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring a DHCP SIP Server on page 450

source-address (SRC Software)

Syntax source-address *source-address*;

Hierarchy Level [edit system services service-deployment]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable the Junos OS to work with the Session and Resource Control (SRC) software.

Options ***source-address***— Local IPv4 address to be used as source address for traffic to the SRC server. The source address restricts traffic within the out-of-band network.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring the Junos OS to Work with SRC Software

source-address-giaddr

Syntax source-address-giaddr;

Hierarchy Level [edit forwarding-options helpers bootp],
[edit forwarding-options helpers bootp interface *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the gateway IP address (giaddr) as the source IP address of the switch for relayed DHCP packets when the switch is used as the DHCP relay agent.

When this statement is entered in the [**edit forwarding-options helpers bootp**] hierarchy, the gateway IP address is configured as the source IP address of the switch for relayed DHCP packets exiting all interfaces on the switch.

When this statement is entered in the [**edit forwarding-options helpers bootp interface *interface-name***] hierarchy, the gateway IP address is configured as the source IP address of the switch for relayed DHCP packets exiting the specified interface of the switch.

The IP address of the interface that the DHCP packet exits on the switch acting as a DHCP relay agent is used as the source IP address for relayed DHCP packets by default.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- DHCP/BOOTP Relay for J-EX Series Switches Overview on page 446

ssh

Syntax
```
ssh {
  protocol-version [v1 v2];
  <connection-limit limit>;
  <rate-limit limit>;
  root-login (allow | deny | deny-password);
}
```

Hierarchy Level [edit system services]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Allow SSH requests from remote systems to the local router or switch.

The remaining statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring SSH Service for Remote Access to the Router or Switch

static-binding

Syntax

```
static-binding mac-address {
  client-identifier (ascii client-id | hexadecimal client-id);
  fixed-address {
    address;
  }
  host client-hostname;
}
```

Hierarchy Level [edit system services **dhcp**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For J-EX Series switches . Set static bindings for DHCP clients. A static binding is a mapping between a fixed IP address and the client's MAC address or client identifier.

Options

mac-address—The MAC address of the client. This is a hardware address that uniquely identifies a client on the network.

fixed-address *address*—Fixed IP address assigned to the client. Typically a client has one address assigned, but you can assign more.

host *client-hostname*—Hostname of the client requesting the DHCP server. The name can include the local domain name. Otherwise, the name is resolved based on the **domain-name** statement.

client-identifier (ascii *client-id* | hexadecimal *client-id*)—Used by the DHCP server to index the database of address bindings. The client identifier is an ASCII string or hexadecimal number and can include a type-value pair as specified in RFC 1700, *Assigned Numbers*. Either a client identifier or the client's MAC address must be configured to uniquely identify the client on the network.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

telnet

Syntax

```
telnet {
  connection-limit limit;
  rate-limit limit;
}
```

Hierarchy Level [edit system services]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Allow Telnet connections from remote systems to the local router or switch.

The remaining statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

tftp

Syntax

```
tftp {
  description text-description;
  interface interface-name {
    broadcast;
    description text-description;
    no-listen;
    server address <logical-system logical-system-name> <routing-instance
      routing-instance-name>;
  }
  server address <logical-system logical-system-name> <routing-instance
    routing-instance-name>;
}
```

Hierarchy Level [edit forwarding-options helpers]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable TFTP request packet forwarding.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring DNS and TFTP Packet Forwarding

traceoptions

Syntax

```
traceoptions {
  file filename <files number> <size size>;
  flag all;
  flag database;
  flag general;
  flag ike;
  flag parse;
  flag policy-manager;
  flag routing-socket;
  flag timer;
}
```

Hierarchy Level

[edit security],
[edit services ipsec-vpn]

Trace options can be configured at either the **[edit security]** or the **[edit services ipsec-vpn]** hierarchy level, but not at both levels.

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure security trace options.

To specify more than one trace option, include multiple **flag** statements. Trace option output is recorded in the **/var/log/kmd** file.

Options

files ***number***—(Optional) Maximum number of trace files. When a trace file (for example, **kmd**) reaches its maximum size, it is renamed **kmd.0**, then **kmd.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 0 files

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB). When a trace file (for example, **kmd**) reaches this size, it is renamed, **kmd.0**, then **kmd.1** and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

Default: 1024 KB

flag—Trace operation to perform. To specify more than one trace operation, include multiple **flag** statements.

- **all**—Trace all security events.
- **database**—Trace database events.
- **general**—Trace general events.
- **ike**—Trace IKE module processing.

- **parse**—Trace configuration processing.
- **policy-manager**—Trace policy manager processing.
- **routing-socket**—Trace routing socket messages.
- **timer**—Trace internal timer events.

Required Privilege Level

admin—To view the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Tracing Operations for Security Services

traceoptions (DHCP Server)

Syntax

```
traceoptions {
  file filename <files number> <match regex> <size size> <world-readable |
    no-world-readable>;
  flag flag;
}
```

Hierarchy Level [edit system services **dhcp**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define tracing operations for DHCP processes for J-EX Series switches.

Options

file *filename*—Name of the file that receives the output of the tracing operation. Enclose the name in quotation marks. All files are placed in the directory **/var/log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

Range: **2** through **1000**

Default: **3** files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **all**—All tracing operations
- **binding**—Trace binding operations
- **config**—Log reading of configuration
- **conflict**—Trace user-detected conflicts for IP addresses
- **event**—Trace important events
- **ifdb**—Trace interface database operations
- **io**— Trace I/O operations
- **lease**—Trace lease operations
- **main**—Trace main loop operations
- **misc**— Trace miscellaneous operations
- **packet**—Trace DHCP packets
- **options**—Trace DHCP options
- **pool**—Trace address pool operations

- **protocol**—Trace protocol operations
- **rtsock**—Trace routing socket operations
- **scope**—Trace scope operations
- **signal**—Trace DHCP signal operations
- **trace**—All tracing operations
- **ui**—Trace user interface operations

match *regex*—(Optional) Refine the output to include lines that contain the regular expression.

- **all**—All tracing operations
- **binding**—Trace binding operations
- **config**— Log reading of configuration
- **conflict**—Trace user-detected conflicts for IP addresses
- **event**—Trace important events
- **ifdb**— Trace interface database operations
- **io**—Trace I/O operations
- **lease**—Trace lease operations
- **main**—Trace main loop operations
- **match *regex***— Refine the output to include lines that contain the regular expression.
- **misc**—Trace miscellaneous operations
- **packet**—Trace DHCP packets
- **options**—Trace DHCP options
- **pool**—Trace address pool operations
- **protocol**—Trace protocol operations
- **rtsock**—Trace routing socket operations
- **scope**—Trace scope operations
- **signal**—Trace DHCP signal operations
- **trace**—All tracing operations
- **ui**—Trace user interface operations

no-world-readable—(Optional) Disable unrestricted file access.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and filename.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10 KB through 1 GB

Default: 128 KB

world-readable—(Optional) Enable unrestricted file access.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Tracing Operations for DHCP Processes
- System Management Configuration Statements

traceoptions (DNS and TFTP Packet Forwarding)

Syntax

```
traceoptions {
  file filename <files number> <match regular-expression> <size bytes> <world-readable |
    no-world-readable>;
  flag flag;
  level level;
  <no-remote-trace>;
}
```

Hierarchy Level [edit forwarding-options **helpers**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure tracing operations for BOOTP, DNS and TFTP packet forwarding.

Default If you do not include this statement, no tracing operations are performed.

Options

file ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks (" "). All files are placed in a file named **fud** in the directory **/var/log**. If you include the **file** statement, you must specify a filename.

files ***number***—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

Range: 2 through 1000

Default: 3 files

flag ***flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **address**—Trace address management events
- **all**—Trace all events
- **bootp**—Trace BOOTP or DHCP services events
- **config**—Trace configuration events
- **domain**—Trace DNS service events
- **ifdb**—Trace interface database operations
- **io**—Trace I/O operations
- **main**—Trace main loop events
- **port**—Trace arbitrary protocol events
- **rtsock**—Trace routing socket operations

- **tftp**—Trace TFTP service events
- **trace**—Trace tracing operations
- **ui**—Trace user interface operations
- **util**—Trace miscellaneous utility operations

match *regular-expression*—(Optional) Refine the output to include lines that contain the regular expression.

no-remote-trace—(Optional) Disable remote tracing globally or for a specific tracing operation.

no-world-readable—(Optional) Restrict file access to the owner.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches this size, it is renamed **trace-file.0**. When the **trace-file** file again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and filename.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 0 bytes through 4,294,967,295 KB

Default: 128 KB

world-readable—(Optional) Enable unrestricted file access.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Tracing BOOTP, DNS, and TFTP Forwarding Operations

web-management

Syntax

```
web-management {
  http {
    interfaces [ interface-names ];
    port port;
  }
  https {
    interfaces [ interface-names ];
    local-certificate name;
    port port;
  }
}
```

Hierarchy Level [edit system services]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure settings for HTTP or HTTPS access. HTTP access allows management of the router or switch using the browser-based J-Web graphical user interface. HTTPS access allows secure management of the router or switch using the J-Web interface. With HTTPS access, communication between the router or switch Web server and your browser is encrypted.

The remaining statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Table 56 on page 396
- *J-Web Interface User Guide*
- **http on page 472**
- **https on page 473**
- **port on page 486**

wins-server

Syntax

```
wins-server {
  address;
}
```

Hierarchy Level [edit system services dhcp],
[edit system services dhcp pool],
[edit system services dhcp static-binding]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For J-EX Series switches . Specify one or more NetBIOS Name Servers. When a DHCP client is added to the network and assigned an IP address, the NetBIOS Name Server manages the Windows Internet Name Service (WINS) database that matches IP addresses (such as **192.168.1.3**) to Windows NetBIOS names (such as **\\Marketing**). List servers in order of preference.

Options ***address***—IPv4 address of the NetBIOS Name Server running WINS. To configure multiple servers, include multiple ***address*** options.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

CHAPTER 38

Operational Mode Commands for System Services

clear system services dhcp binding

Syntax clear system services dhcp binding
<*address*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description (J-EX Series switches) Remove obsolete IP address bindings on a Dynamic Host Configuration Protocol (DHCP) server and return them to the IP address pool.

Options *address*—(Optional) Remove a specific IP address binding and return it to the address pool.

Required Privilege Level view and system

Related Documentation
- **show system services dhcp binding on page 517**

List of Sample Output **clear system services dhcp binding on page 510**

Output Fields When you enter this command, you are provided feedback on the status of your request.

clear system services dhcp binding

```
user@host> clear system services dhcp binding
```

clear system services dhcp conflict

Syntax clear system services dhcp conflict
<*address*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description (J-EX Series switches) Remove IP addresses from the Dynamic Host Configuration Protocol (DHCP) server conflict list and return them to the IP address pool.

Options *address*—(Optional) Remove a specific IP address from the conflict list and return it to the address pool.

Required Privilege Level view and system

Related Documentation
- **show system services dhcp conflict on page 519**

List of Sample Output **clear system services dhcp conflict on page 511**

Output Fields When you enter this command, you are provided feedback on the status of your request.

clear system services dhcp conflict

```
user@host> clear system services dhcp conflict
```

clear system services dhcp statistics

Syntax clear system services dhcp statistics

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description (J-EX Series switches) Clear Dynamic Host Configuration Protocol (DHCP) server statistics.

Options This command has no options.

Required Privilege Level view and system

Related Documentation

- **show system services dhcp statistics on page 524**

List of Sample Output **clear system services dhcp statistics on page 512**

Output Fields When you enter this command, you are provided feedback on the status of your request.

clear system services dhcp statistics

```
user@host> clear system services dhcp statistics
```

request ipsec switch

Syntax request ipsec switch (interface <es-*fpc/pic/port*> | security-associations <*sa-name*>)

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description (Encryption interface on J-EX Series switches) Manually switch from the primary to the backup encryption services interface, or switch from the primary to the backup IP Security (IPsec) tunnel.

Options interface <es-*fpc/pic/port*>—Switch to the backup encryption interface.

security-associations <*sa-name*>—Switch to the backup tunnel.

Required Privilege Level view

Related Documentation
- show ipsec redundancy

List of Sample Output **request ipsec switch on page 513**

Output Fields When you enter this command, you are provided feedback on the status of your request.

request ipsec switch

```
user@host> request ipsec switch security-associations sa-private
```

request security certificate (signed)

Syntax request security certificate enroll filename *filename* subject *subject* alternative-subject *alternative-subject* certification-authority *certification-authority* encoding (binary | pem) key-file *key-file* domain-name *domain-name*

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description (Encryption interface on J-EX Series switches) Obtain a signed certificate from a certificate authority (CA). The signed certificate validates the CA and the owner of the certificate. The results are saved in a specified file to the **/var/etc/ikecert** directory.

Options filename *filename*—File that stores the certificate.

subject *subject*—Distinguished name (**dn**), which consists of a set of components—for example, an organization (**o**), an organization unit (**ou**), a country (**c**), and a locality (**l**).

alternative-subject *alternative-subject*—Tunnel source address.

certification-authority *certification-authority*—Name of the certificate authority profile in the configuration.

encoding (binary | pem)—File format used for the certificate. The format can be a binary file or privacy-enhanced mail (PEM), an ASCII base64-encoded format. The default format is binary.

key-file *key-file*—File containing a local private key.

domain-name *domain-name*—Fully qualified domain name.

Required Privilege Level maintenance

List of Sample Output **request security certificate (signed) on page 514**

Output Fields When you enter this command, you are provided feedback on the status of your request.

request security certificate (signed)

```
user@host> request security certificate enroll filename host.crt subject c=uk,o=london
alternative-subject 10.50.1.4 certification-authority verisign  key-file host-1.prv domain-name
host.juniper.net
CA name: juniper.net CA file: ca_verisign
local pub/private key pair: host.prv
subject: c=uk,o=london domain name: host.juniper.net
alternative subject: 10.50.1.4
Encoding: binary
Certificate enrollment has started. To view the status of your enrollment, check
 the key management process (kmd) log file at /var/log/kmd. <--------------
```

request security key-pair

Syntax
request security key-pair *filename*
<size *key-size*>
<type (rsa | dsa)>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
(Encryption interface on J-EX Series switches) Generate a public and private key pair for a digital certificate.

Options
filename—Name of a file in which to store the key pair.

size *key-size*—(Optional) Key size, in bits. The key size can be **512**, **1024**, or **2048**. The default value is **1024**.

type—(Optional) Algorithm used to encrypt the key:

- **rsa**—RSA algorithm. This is the default.
- **dsa**—Digital signature algorithm with Secure Hash Algorithm (SHA).

Required Privilege Level
maintenance

List of Sample Output
request security key-pair on page 515

Output Fields
When you enter this command, you are provided feedback on the status of your request.

request security key-pair
`user@host>` **request security key-pair security-key-file**

request security certificate (unsigned)

Syntax request security certificate enroll filename *filename* ca-file *ca-file* ca-name *ca-name* encoding (binary | perm) url *url*

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description (Encryption interface on J-EX Series switches) Obtain a certificate from a certificate authority (CA). The results are saved in a specified file to the **/var/etc/ikecert** directory.

Options

filename *filename*—File that stores the public key certificate.

ca-file *ca-file*—Name of the certificate authority profile in the configuration.

ca-name *ca-name*—Name of the certificate authority.

encoding (binary | pem)—File format used for the certificate. The format can be a binary file or privacy-enhanced mail (PEM), an ASCII base64-encoded format. The default value is **binary**.

url *url*—Certificate authority URL.

Required Privilege Level maintenance

List of Sample Output **request security certificate (unsigned) on page 516**

Output Fields When you enter this command, you are provided feedback on the status of your request.

request security certificate (unsigned)

```
user@host> request security certificate enroll filename ca_verisign ca-file verisign ca-name
juniper.net urlxyzcompany URL
http://<verisign ca-name xyzcompany url>/cgi-bin/pkiclient.exe CA name: juniper.net
 CA file: verisign Encoding: binary
Certificate enrollment has started. To view the status of your enrollment, check
 the key management process (kmd) log file at /var/log/kmd. <--------------
```

show system services dhcp binding

Syntax show system services dhcp binding
<detail>
<*address*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display Dynamic Host Configuration Protocol (DHCP) server client binding information.

Options none—Display brief information about all active client bindings.

detail—(Optional) Display detailed information about all active client bindings.

address—(Optional) Display detailed client binding information for the specified IP address only.

Required Privilege Level view and system

Related Documentation
- **clear system services dhcp binding on page 510**

List of Sample Output
show system services dhcp binding on page 518
show system services dhcp binding address on page 518
show system services dhcp binding address detail on page 518

Output Fields Table 62 on page 517 describes the output fields for the **show system services dhcp binding** command. Output fields are listed in the approximate order in which they appear.

Table 62: show system services dhcp binding Output Fields

Field Name	Field Description	Level of Output
Allocated address	List of IP addresses the DHCP server has assigned to clients.	All levels
MAC address	Corresponding media access control (MAC) hardware address of the client.	All levels
Client identifier	(***address*** option only) Client's unique identifier (represented by an ASCII string or hexadecimal digits). This identifier is used by the DHCP server to index its database of address bindings.	All levels
Binding Type	Type of binding assigned to the client. DHCP servers can assign a dynamic binding from a pool of IP addresses or a static binding to one or more specific IP addresses.	All levels
Lease Expires at	Time the lease expires or **never** for leases that do not expire.	All levels
Lease Obtained at	(***address*** option only) Time the client obtained the lease from the DHCP server.	**detail**
State	Status of the binding. Bindings can be active or expired.	**detail**

Table 62: show system services dhcp binding Output Fields *(continued)*

Field Name	Field Description	Level of Output
Pool	Address pool that contains the IP address assigned to the client.	**detail**
Request received on	Interface on which the DHCP message exchange occurs. The IP address pool is configured based on the interface's IP address. If a relay agent is used, its IP address is also displayed.	**detail**
DHCP options	User-defined options created for the DHCP server. If no options have been defined, this field is blank.	**detail**

show system services dhcp binding

```
user@host> show system services dhcp binding

Allocated address   MAC address        Binding Type   Lease expires at
192.168.1.2         00:a0:12:00:12:ab  static         never
192.168.1.3         00:a0:12:00:13:02  dynamic        2004-05-03 13:01:42 PDT
```

show system services dhcp binding address

```
user@host> show system services dhcp binding 192.168.1.3

DHCP binding information:
Allocated address: 192.168.1.3
Mac address: 00:a0:12:00:12:ab
Client identifier
61 63 65 64 2d 30 30 3a 61 30 3a 31 32 3a 30 30aced-00:a0:12:00
3a 31 33 3a 30 32:13:02

Lease information:
    Binding Type dynamic
    Obtained at 2004-05-02 13:01:42 PDT
    Expires at 2004-05-03 13:01:42 PDT
```

show system services dhcp binding address detail

```
user@host> show system services dhcp binding 192.168.1.3 detail
DHCP binding information:
Allocated address        192.168.1.3
MAC address  00:a0:12:00:12:ab
    Pool              192.168.1.0/24
Request received on fe-0/0/0, relayed by 192.168.4.254

Lease information:
    Type              DHCP
    Obtained at       2004-05-02 13:01:42 PDT
    Expires at        2004-05-03 13:01:42 PDT
State active

DHCP options:
    Name: name-server, Value: { 6.6.6.6, 6.6.6.7 }
    Name: domain-name, Value: mydomain.tld
    Code: 19, Type: flag, Value: off
    Code: 40, Type: string, Value: domain.tld
    Code: 32, Type: ip-address, Value: 3.3.3.33
```

show system services dhcp conflict

Syntax show system services dhcp conflict

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description (J-EX Series switches) Display Dynamic Host Configuration Protocol (DHCP) client-detected conflicts for IP addresses. When a conflict is detected, the DHCP server removes the address from the address pool.

Options This command has no options.

Required Privilege Level view and system

Related Documentation
- **clear system services dhcp conflict on page 511**

List of Sample Output **show system services dhcp conflict on page 519**

Output Fields Table 63 on page 519 describes the output fields for the **show system services dhcp conflict** command. Output fields are listed in the approximate order in which they appear.

Table 63: show system services dhcp conflict Output Fields

Field Name	Field Description
Detection time	Date and time the client detected the conflict.
Detection method	How the conflict was detected.
Address	IP address where the conflict occurs. The addresses in the conflicts list remain excluded from the pool until you use a **clear system services dhcp conflict** command to manually clear the list.

show system services dhcp conflict

```
user@host> show system services dhcp conflict

Detection time           Detection method  Address
2004-08-03 19:04:00 PDT  ARP               3.3.3.5
2004-08-04 04:23:12 PDT  Ping              4.4.4.8
2004-08-05 21:06:44 PDT  Client            3.3.3.10
```

show system services dhcp global

Syntax show system services dhcp global

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description (J-EX Series switches) Display Dynamic Host Configuration Protocol (DHCP) global configuration options. Global options apply to all scopes and clients served by the DHCP server. Global options are overridden if specified otherwise in scope or client options. Scope options apply to specific subnets or ranges of addresses. Client options apply to specific clients.

Options This command has no options.

Required Privilege Level view and system

List of Sample Output **show system services dhcp global on page 521**

Output Fields Table 64 on page 520 describes the output fields for the **show system services dhcp global** command. Output fields are listed in the approximate order in which they appear.

Table 64: show system services dhcp global Output Fields

Field Name	Field Description
BOOTP lease length	Length of lease time assigned to BOOTP clients.
Default lease time	Lease time assigned to clients that do not request a specific lease time.
Minimum lease time	Minimum time a client retains an IP address lease on the server.
Maximum lease time	Maximum time a client can retain an IP address lease on the server.
DHCP options	User-defined options created for the DHCP server. If no options have been defined, this field is blank.

show system services dhcp global

```
user@host> show system services dhcp global

Global settings:
    BOOTP lease length         infinite

DHCP lease times:
    Default lease time         1 hour
    Minimum lease time         2 hours
    Maximum lease time         infinite

DHCP options:
    Name: name-server, Value: { 6.6.6.6, 6.6.6.7 }
    Name: domain-name, Value: mydomain.tld
    Code: 19, Type: flag, Value: off
    Code: 40, Type: string, Value: domain.tld
    Code: 32, Type: ip-address, Value: 3.3.3.33
```

show system services dhcp pool

Syntax show system services dhcp pool
<detail>
<*subnet-address*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description (J-EX Series switches) Display Dynamic Host Configuration Protocol (DHCP) server IP address pools.

Options none—Display brief information about all IP address pools.

detail—(Optional) Display detailed information.

subnet-address—(Optional) Display information for the specified subnet address.

Required Privilege Level view and system

List of Sample Output **show system services dhcp pool on page 523**
show system services dhcp pool subnet-address on page 523
show system services dhcp pool subnet-address detail on page 523

Output Fields Table 65 on page 522 describes the output fields for the **show system services dhcp pool** command. Output fields are listed in the approximate order in which they appear.

Table 65: show system services dhcp pool Output Fields

Field Name	Field Description	Level of Output
Pool name	Subnet on which the IP address pool is defined.	None specified
Low address	Lowest address in the IP address pool.	None specified
High address	Highest address in the IP address pool.	None specified
Excluded addresses	Addresses excluded from the address pool.	None specified
Subnet	(***subnet-address*** option only) Subnet to which the specified address pool belongs.	None specified
Address range	(***subnet-address*** option only) Range of IP addresses in the address pool.	None specified
Addresses assigned	Number of IP addresses in the pool that are assigned to DHCP clients and the total number of IP addresses in the pool.	**detail**
Active	Number of assigned IP addresses in the pool that are active.	**detail**
Excluded	Number of assigned IP addresses in the pool that are excluded.	**detail**
Default lease time	Lease time assigned to clients that do not request a specific lease time.	**detail**

Table 65: show system services dhcp pool Output Fields *(continued)*

Field Name	Field Description	Level of Output
Minimum lease time	Minimum time a client can retain an IP address lease on the server.	**detail**
Maximum lease time	Maximum time a client can retain an IP address lease on the server.	**detail**
DHCP options	User-defined options created for the DHCP server. If no options have been defined, this field is blank.	**detail**

show system services dhcp pool

```
user@host> show system services dhcp pool

Pool name    Low address   High address   Excluded addresses
3.3.3.0/24   3.3.3.2       3.3.3.254      3.3.3.1
```

show system services dhcp pool subnet-address

```
user@host> show system services dhcp pool 3.3.3.0/24

Pool information:
    Subnet                      3.3.3.0/24
    Address range               3.3.3.2 - 3.3.3.254
    Addresses assigned          2/253
```

show system services dhcp pool subnet-address detail

```
user@host> show system services dhcp pool 3.3.3.0/24 detail

Pool information:
    Subnet                      3.3.3.0/24
    Address range               3.3.3.2 - 3.3.3.254
    Addresses assigned          2/253
      Active: 1, Excluded: 1

DHCP lease times:
    Default lease time          1 hour
    Minimum lease time          2 hours
    Maximum lease time          infinite

DHCP options:
    Name: name-server, Value: { 6.6.6.6, 6.6.6.7 }
    Name: domain-name, Value: mydomain.tld
    Name: router, Value: { 3.3.3.1 }
    Name: server-identifier, Value: 3.3.3.1
    Code: 19, Type: flag, Value: off
    Code: 40, Type: string, Value: domain.tld
    Code: 32, Type: ip-address, Value: 3.3.3.333.3.3.254 3.3.3.1
```

show system services dhcp statistics

Syntax show system services dhcp statistics

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description (J-EX Series switches) Display Dynamic Host Configuration Protocol (DHCP) server statistics.

Options This command has no options.

Required Privilege Level view and system

Related Documentation
- **clear system services dhcp statistics on page 512**

List of Sample Output **show system services dhcp statistics on page 525**

Output Fields Table 66 on page 524 describes the output fields for the **show system services dhcp statistics** command. Output fields are listed in the approximate order in which they appear.

Table 66: show system services dhcp statistics Output Fields

Field Name	Field Description
Default lease time	Lease time assigned to clients that do not request a specific lease time.
Minimum lease time	Minimum time a client can retain an IP address lease on the server.
Maximum lease time	Maximum time a client can retain an IP address lease on the server.
Packets dropped	Total number of packets dropped and number of packets dropped because of: • Invalid hardware address • Invalid opcode • Invalid server address • No available address • No interface match • No routing instance match • No valid local addresses • Packet too short • Read error • Send error

Table 66: show system services dhcp statistics Output Fields *(continued)*

Field Name	Field Description
Messages received	Number of the following message types sent from DHCP clients and received by the DHCP server: • BOOTREQUEST • DHCPDECLINE • DHCPDISCOVER • DHCPINFORM • DHCPRELEASE • DHCPREQUEST
Messages sent	Number of the following message types sent from the DHCP server to DHCP clients: • BOOTREPLY • DHCPACK • DHCPOFFER • DHCPNAK

show system services dhcp statistics

```
user@host> show system services dhcp statistics

DHCP lease times:
    Default lease time         1 hour
    Minimum lease time         2 hours
    Maximum lease time         infinite

Packets dropped:
    Total                      0
    Bad hardware address       0
    Bad opcode                 0
    Invalid server address     0
    No available addresses     0
    No interface match         0
    No routing instance match  0
    No valid local address     0
    Packet too short           0
    Read error                 0
    Send error                 0

Messages received:
    BOOTREQUEST                0
    DHCPDECLINE                0
    DHCPDISCOVER               0
    DHCPINFORM                 0
    DHCPRELEASE                0
    DHCPREQUEST                0

Messages sent:
    BOOTREPLY                  0
    DHCPACK                    0
    DHCPOFFER                  0
    DHCPNAK                    0
```

show system services service-deployment

Syntax show system services service-deployment

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display information about a Session and Resource Control (SRC) client.

Options This command has no options.

Required Privilege Level view and system

List of Sample Output **show system services service-deployment on page 526**

show system services service-deployment

```
user@host> show system services service-deployment
Connected to 192.4.4.4 port 10288 since 2004-05-03 11:04:34 PDT Keepalive settings:
 Interval 15 seconds Keepalives sent: 750 Notifications sent: 0 Last update from
 peer: 00:00:06 ago
```

ssh

Syntax

ssh *host*
<bypass-routing>
<inet | inet6>
<interface *interface-name*>
<logical-system *logical-system-name*>
<routing-instance *routing-instance-name*>
<source *address*>
<v1 | v2>

Syntax (J-EX Series Switch)

ssh *host*
<bypass-routing>
<inet | inet6>
<interface *interface-name*>
<routing-instance *routing-instance-name*>
<source *address*>
<v1 | v2>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Use the SSH program to open a connection between a local router or switch and a remote system and execute commands on the remote system. You can issue the **ssh** command from the Junos OS CLI to log in to a remote system or from a remote system to log in to the local router or switch. When executing this command, you include one or more CLI commands by enclosing them in quotation marks and separating the commands with semicolons:

ssh *address* '*cli-command1* ; *cli-command2* '

Options

host—Name or address of the remote system.

bypass-routing—(Optional) Bypass the normal routing tables and send ping requests directly to a system on an attached network. If the system is not on a directly attached network, an error is returned. Use this option to ping a local system through an interface that has no route through it.

inet | inet6—(Optional) Create an IPv4 or IPv6 connection, respectively.

interface *interface-name*—(Optional) Interface name for the SSH session. (This option does not work when **default-address-selection** is configured at the **[edit system]** hierarchy level, because this configuration uses the loopback interface as the source address for all locally generated IP packets.)

logical-system *logical-system-name*—(Optional) Name of a particular logical system for the SSH attempt.

routing-instance *routing-instance-name*—(Optional) Name of the routing instance for the SSH attempt.

source *address*—(Optional) Source address of the SSH connection.

v1 | v2—(Optional) Use SSH version 1 or 2, respectively, when connecting to a remote host.

Additional Information

To configure an SSH (version 1) key for your user account, include the **authentication ssh-rsa** statement at the **[edit system login user *user-name*]** hierarchy level. To configure an SSH (version 2) key for your user account, include the **authentication dsa-rsa** statement at the **[edit system login user *user-name*]** hierarchy level. For details, see the *Junos OS System Basics Configuration Guide*.

You can limit the number of times a user can attempt to enter a password while logging in through SSH. To specify the number of times a user can attempt to enter a password to log in through SSH, include the **retry-options** statement at the **[edit system login]** hierarchy level. For details, see the *Junos OS System Basics Configuration Guide*.

Required Privilege Level

network

List of Sample Output

ssh on page 528

Output Fields

When you enter this command, you are provided feedback on the status of your request.

ssh

```
user@host> ssh cree
Host key not found from the list of known hosts.
Are you sure you want to continue connecting (yes/no)? yes

Host ?cree' added to the list of known hosts.
boojun@cree's password:
Last login: Sun Jun 21 10:43:42 1998 from junos-router
% ...
```

telnet

Syntax

telnet *host*
<8bit>
<bypass-routing>
<inet | inet6>
<interface *interface-name*>
<logical-system *logical-system-name*>
<no-resolve>
<port *port-number*>
<routing-instance *routing-instance-name*>
<source *source-address*>

Syntax (J-EX Series Switch)

telnet *host*
<8bit>
<bypass-routing>
<inet | inet6>
<interface *interface-name*>
<no-resolve>
<port *port-number*>
<routing-instance *routing-instance-name*>
<source *source-address*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Open a telnet session to a remote system. Type Ctrl+] to escape from the telnet session to the telnet command level, and then type **quit** to exit from telnet.

Options

host—Name or address of the remote system.

8bit—(Optional) Use an 8-bit data path.

bypass-routing—(Optional) Bypass the normal routing tables and send ping requests directly to a system on an attached network. If the system is not on a directly attached network, an error is returned. Use this option to ping a local system through an interface that has no route through it.

inet | inet6—(Optional) Open an IPv4 or IPv6 session, respectively.

interface *interface-name*—(Optional) Interface name for the telnet session. (This option does not work when **default-address-selection** is configured at the **[edit system]** hierarchy level, because this configuration uses the loopback interface as the source address for all locally generated IP packets.)

logical-system *logical-system-name*—(Optional) Name of a particular logical system for the telnet attempt.

no-resolve—(Optional) Do not attempt to determine the hostname that corresponds to the IP address.

port *port-number*—(Optional) Port number or service name on the remote system.

routing-instance *routing-instance-name*—(Optional) Name of the routing instance for the telnet attempt.

source *source-address*—(Optional) Source address of the telnet connection.

Additional Information You can limit the number of times a user can attempt to enter a password while logging in through telnet. To specify the number of times a user can attempt to enter a password to log in through telnet, include the **retry-options** statement at the [**edit system login**] hierarchy level. For details, see the *Junos OS System Basics Configuration Guide*.

Required Privilege Level network

List of Sample Output **telnet on page 530**

Output Fields When you enter this command, you are provided feedback on the status of your request.

telnet

```
user@host> telnet 192.154.1.254
Trying 192.154.169.254...
Connected to level5.company.net.
Escape character is '^]'.
ttypa
login:
```

PART 10

Junos OS for J-EX Series Switches System Monitoring

CHAPTER 39

System Monitoring Overview

Understanding Alarm Types and Severity Levels on J-EX Series Switches

Before monitoring alarms on the switch, become familiar with the terms defined in Table 67 on page 533.

Table 67: Alarm Terms

Term	Definition
alarm	Signal alerting you to conditions that might prevent normal operation. On a switch, the alarm signal is the yellow **ALARM** LED lit on the front of the chassis.
alarm condition	Failure event that triggers an alarm.
alarm severity	Seriousness of the alarm. The level of severity can be either major (red) or minor (yellow).
chassis alarm	Predefined alarm triggered by a physical condition on the switch such as a power supply failure, excessive component temperature, or media failure.
system alarm	Predefined alarm triggered by a missing rescue configuration or failure to install a license for a licensed software feature.

Alarm Types

The switch supports these alarms:

- Chassis alarms indicate a failure on the switch or one of its components. Chassis alarms are preset and cannot be modified.
- System alarms indicate a missing rescue configuration. System alarms are preset and cannot be modified, although you can configure them to appear automatically in the J-Web interface display or CLI display.

Alarm Severity Levels

Alarms on a J-EX Series Switches have two severity levels:

- Major (red)—Indicates a critical situation on the switch that has resulted from one of the following conditions. A red alarm condition requires immediate action.
  - One or more hardware components have failed.
  - One or more hardware components have exceeded temperature thresholds.
  - An alarm condition configured on an interface has triggered a critical warning.

  NOTE: When you connect power to a J-EX4200 switch, the Alarm (ALM) LED lights red to indicate that the network link is disconnected. This behavior is normal. Plugging an active Ethernet cable into the management (MGMT) port on the switch completes the network link and turns off the ALM LED.

  Connecting the switch to a dedicated management console instead of a network does not affect the ALM LED. The LED remains red until the switch is connected to a network.

- Minor (yellow or amber)—Indicates a noncritical condition on the switch that, if left unchecked, might cause an interruption in service or degradation in performance. A yellow alarm condition requires monitoring or maintenance.

  A missing rescue configuration generates a yellow system alarm. To set the rescue configuration, see “Setting or Deleting the Rescue Configuration (CLI Procedure)” on page 344.

Related Documentation

- Checking Active Alarms with the J-Web Interface on page 544
- Dashboard for J-EX Series Switches on page 534

Dashboard for J-EX Series Switches

When you log in to the J-Web user interface, the dashboard for the J-EX Series switch appears. Use the dashboard to view system information.

The dashboard comprises four panels and a graphical chassis viewer. You can click **Preferences** to choose which panels are to be displayed and set the refresh interval for chassis viewer information. Click **OK** to save your preference changes and return to the dashboard or click **Cancel** to return to the dashboard without saving changes.

NOTE: You can drag and drop the various panels to different locations in the J-Web window.

This topic describes:

System Information Panel

Table 68: System Information

Field	Description
System name	Indicates the local name of the J-EX Series switch.
Device model	Indicates the model of the J-EX Series switch. NOTE: For a J-EX8208 switch or a J-EX8216 switch, the **Device model** information changes with respect to the selected line card, the Switch Fabric and Routing Engine (SRE) module in a J-EX8208 switch, or the Routing Engine (RE) module in a J-EX8216 switch.
Inventory details	Indicates the following: • For J-EX4200 switches not configured as Virtual Chassis, the value in **Inventory** is always 1 FPC. FPC is a legacy term for a slot in a large router chassis; here, it simply refers to the single switch. • For a J-EX4200 switch configured as a Virtual Chassis, the value in **Inventory** is displayed as 1–10 FPC, with the number corresponding to the number of member switches. • For a J-EX8208 switch, the values in **Inventory** are displayed as 1–3 CB and 0–8 FPC. Control board (CB) refers to SRE and SF modules. FPC refers to line cards. • For a J-EX8216 switch, the values in **Inventory** are displayed as 1-2 CB and 0-16 FPC. Control board (CB) refers to RE modules. FPC refers to line cards.
Junos image	Indicates the version of the Junos OS image.
Boot image	Indicates the version of the boot image that is used.
Device uptime	Indicates the time since the last reboot.
Last configured time	Indicates the time when the switch was last configured.

Health Status Panel

Table 69: Health Status

Field	Description
J-EX4200 Switches	
Memory util.	Indicates the memory used in the Routing Engine. In a Virtual Chassis configuration, the memory utilization value of the master Routing Engine is displayed.
Flash	Indicates the usage and capacity of internal flash memory and any external USB flash drive.

Table 69: Health Status *(continued)*

Field	Description
Temp.	Indicates the chassis temperature status. Temperatures in the dashboard are listed in Celsius and the corresponding Fahrenheit values.
CPU load	Indicates the average CPU usage over 15 minutes.
Fan status	Indicates the fan status of the switch. The possible values are **OK**, **Failed**, and **Absent**.
J-EX8208 Switches	
Memory util.	Indicates the memory used in the Routing Engine. If there are two Routing Engines, the memory utilization value of the master is displayed.
CPU load	Indicates the average CPU usage over 15 minutes.
Flash	Indicates the usage and capacity of internal flash memory and any external USB flash drive.
J-EX8216 Switches	
Memory util.	Indicates the memory used in the Routing Engine. If there are two Routing Engines, the memory utilization value of the master is displayed.
CPU load	Indicates the average CPU usage over 15 minutes.
Flash	Indicates the usage and capacity of internal flash memory and any external USB flash drive.

Capacity Utilization Panel

Table 70: Capacity Utilization

Field	Description
Number of active ports	Indicates the number of active ports in the switch.
Total number of ports	Indicates the number of ports in the switch.
Used-up MAC-Table entries	Indicates the number of MAC-Table entries.
Supported MAC-Table entries	Indicates the maximum number of MAC-Table entries permitted.
Number of VLANs configured	Indicates the number of configured VLANs.
Number of VLANs supported	Indicates the maximum number of VLANs that are supported.

Alarms Panel

Displays information about the last five alarms raised in the system. For example, if there are 5 major alarms, then details for all 5 major alarms are displayed. If there are 4 major alarms and 3 minor alarms, then details of the 4 major alarms and 1 minor alarm are displayed. Major alarms are displayed in red and minor alarms are displayed in yellow.

NOTE: When you connect power to a J-EX4200 switch, the Alarm (ALM) LED lights red to indicate that the network link is disconnected. This behavior is normal. Plugging an active Ethernet cable into the management (MGMT) port on the switch completes the network link and turns off the ALM LED.

Connecting the switch to a dedicated management console instead of a network does not affect the ALM LED. The LED remains red until the switch is connected to a network.

Chassis Viewer

You can click the **Rear View** button to see the back of the chassis image. Click **Front View** to see the front of the image. In a Virtual Chassis configuration, the **Rear View** button is disabled if the switch is not selected.

- Table 71 on page 537—Describes the chassis viewer for J-EX4200 switches.
- Table 72 on page 538—Describes the chassis viewer for J-EX8208 switches.
- Table 73 on page 540—Describes the chassis viewer for J-EX8216 switches.

Table 71: Chassis Viewer for J-EX4200 Switches

Field	Description
Front View	
Interface status	In the image, the colors listed below denote the interface status: • Green—Interface is up and operational. • Yellow—Interface is up but is nonoperational. • Gray—Interface is down and nonoperational. Hover the mouse pointer over the interface (port) to view more information. For a Virtual Chassis configuration, select the switch to view the interface status. If an SFP+ uplink module is installed in the switch, hover the mouse pointer over the port icon to display whether the module is configured to operate in 1G mode or 10G mode. If the module is configured to operate in 1G mode, the tool tip information is displayed for all 4 ports. If the module is configured to operate in 10G mode, the tool tip information is displayed only for 2 ports. For SFP and SFP+ ports, the interfaces appear dimmed if no transceiver is inserted. The chassis viewer displays “Transceiver not plugged-in” when you hover the mouse pointer over the port icon.
LCD panel	LCD panel configured for the LEDs on the ports. Hover the mouse pointer over the icon to view the current character display.
Rear View of the J-EX4200 Switch	
Fan tray	Hover the mouse pointer over the fan tray icon to display Name, Status, and Description information. For a Virtual Chassis, the status of the fans of the selected member switch is displayed.

Table 71: Chassis Viewer for J-EX4200 Switches *(continued)*

Field	Description
Virtual Chassis port	Displayed only when switches are configured as a Virtual Chassis. The colors listed below denote the Virtual Chassis port (VCP) status: • Green—VCP is up and operational. • Yellow—VCP is up but is nonoperational. • Gray—VCP is down and nonoperational.
USB port	Indicates the USB port for the switch. NOTE: We recommend you use USB flash drives purchased from Dell for your J-EX Series switch.
Management (**me0**) port	The management port is used to connect the switch to a management device for out-of-band management.
Console port	The console port is used to connect the switch to a management console or to a console server. (You might do this for initial switch configuration.)
Power supplies	Hover the mouse pointer over the power supply icons to display Name, Status, and Description information.

Table 72: Chassis Viewer for J-EX8208 Switches

Field	Description
Front View	
Interface status	In the image, click any line card, SRE module, or SF module to view the front view of the selected component. The colors listed below denote the interface status: • Green—Interface is up and operational. • Yellow—Interface is up but is nonoperational. • Gray—Interface is down and nonoperational. Hover the mouse pointer over the interface (port) to view more information. You can view status for the following ports on the SRE module: • USB port—Indicates the USB port for the switch. NOTE: We recommend you use USB flash drives purchased from Dell for your J-EX Series switch. • Auxiliary port—This port is not enabled on the switch. It is reserved for future use. • Management (**me0**) port—The management port is used to connect the switch to a management device for out-of-band management. • Console port—The console port is used to connect the switch to a management console or to a console server. (You might do this for initial switch configuration.) Because the SF module has no ports, no status information is displayed.

Table 72: Chassis Viewer for J-EX8208 Switches *(continued)*

Field	Description
Slot numbers	Slots on the switch are labeled, from the top of the switch down: • 0–3 (line cards) • SRE0, SF, SRE1 (SRE and SF modules) • 4–7 (line cards)
Temperature	The active slots contain a gray temperature icon. Hover the mouse pointer over the icon to display temperature information for the slot.
Fan status	Hover the mouse pointer over the fan tray icon to display Name, Status, and Description information.
Power supplies	Hover the mouse pointer over the power supply icons to display Name, Status, and Description information.
LCD panel	LCD panel configured for the LEDs on the ports. Hover the mouse pointer over the icon to view the current character display.
Rear View	The J-EX8208 switch does not have any components on the rear of the chassis.

Table 73: Chassis Viewer for J-EX8216 Switches

Field	Description
Front View	
Interface status	In the image, click any line card or RE module to view the front view of the selected component. The colors listed below denote the interface status: • Green—Interface is up and operational. • Yellow—Interface is up but is nonoperational. • Gray—Interface is down and nonoperational. Hover the mouse pointer over the interface (port) to view more information. You can view status for the following ports on the RE module: • USB port—Indicates the USB port for the switch. NOTE: We recommend you use USB flash drives purchased from Dell for your J-EX Series switch. • Auxiliary port—This port is not enabled on the switch. It is reserved for future use. • Management (**me0**) port—The management port is used to connect the switch to a management device for out-of-band management. • Console port—The console port is used to connect the switch to a management console or to a console server. (You might do this for initial switch configuration.)
Slot numbers	Slots on the switch are labeled, from the top of the switch down: • RE0 (RE module) • RE1 (RE module) • 0–15 (line cards)
Temperature	The active slots contain a gray temperature icon. Hover the mouse pointer over the icon to display temperature information for the slot.
Fan status	Hover the mouse pointer over the fan tray icon to display consolidated fan information.
Power supplies	Hover the mouse pointer over the power supply icons to display Name, Status, and Description information.
LCD panel	LCD panel configured for the LEDs on the ports. Hover the mouse pointer over the icon to view the current character display.
Rear View	
SF modules	Hover the mouse pointer over the SF module icons in their respective slots to display information. Slots are numbered SF7–SF0, from left to right.

Related Documentation

- J-Web User Interface for J-EX Series Switches Overview on page 129
- Checking Active Alarms with the J-Web Interface on page 544
- J-EX4200 Switches Hardware Overview on page 25
- J-EX8208 Switch Hardware Overview on page 27
- J-EX8216 Switch Hardware Overview on page 30

CHAPTER 40

Administering and Monitoring System Functions

Monitoring System Log Messages

Purpose Use the monitoring functionality to filter and view system log messages for J-EX Series switches.

Action To view events in the J-Web interface, select **Monitor > Events and Alarms > View Events**.

Apply a filter or a combination of filters to view messages. You can use filters to display relevant events. Table 74 on page 541 describes the different filters, their functions, and the associated actions.

To view events in the CLI, enter the following command:

show log

Table 74: Filtering System Log Messages

Field	Function	Your Action
System Log File	Specifies the name of a system log file for which you want to display the recorded events. Lists the names of all the system log files that you configure. By default, a log file, **messages**, is included in the **/var/log/** directory.	To specify events recorded in a particular file, select the system log filename from the list—for example, **messages**. Select **Include archived files** to include archived files in the search.

Table 74: Filtering System Log Messages *(continued)*

Field	Function	Your Action
Process	Specifies the name of the process generating the events you want to display. To view all the processes running on your system, enter the CLI command **show system processes**. For more information about processes, see the *Junos OS Installation and Upgrade Guide* at http://www.juniper.net/techpubs/software/junos/.	To specify events generated by a process, type the name of the process. For example, type **mgd** to list all messages generated by the management process.
Date From To	Specifies the time period in which the events you want displayed are generated. Displays a calendar that allows you to select the year, month, day, and time. It also allows you to select the local time. By default, the messages generated in the last hour are displayed. End Time shows the current time and Start Time shows the time one hour before End Time.	To specify the time period: • Click the **Calendar** icon and select the year, month, and date—for example, **02/10/2007**. • Click the **Calendar** icon and select the year, month, and date—for example, **02/10/2007**. • Click to select the time in hours, minutes, and seconds.
Event ID	Specifies the event ID for which you want to display the messages. Allows you to type part of the ID and completes the remainder automatically. An event ID, also known as a system log message code, uniquely identifies a system log message. It begins with a prefix that indicates the generating software process or library.	To specify events with a specific ID, type the partial or complete ID—for example, **TFTPD_AF_ERR**.
Description	Specifies text from the description of events that you want to display. Allows you to use regular expressions to match text from the event description. NOTE: Regular expression matching is case-sensitive.	To specify events with a specific description, type a text string from the description with regular expression. For example, type **^Initial*** to display all messages with lines beginning with the term *Initial*.
Search	Applies the specified filter and displays the matching messages.	To apply the filter and display messages, click **Search**.

Meaning Table 75 on page 543 describes the Event Summary fields.

NOTE: By default, the View Events page in the J-Web interface displays the most recent 25 events, with severity levels highlighted in different colors. After you specify the filters, Event Summary displays the events matching the specified filters. Click the **First**, **Next**, **Prev**, and **Last** links to navigate through messages.

Table 75: Viewing System Log Messages

Field	Function	Additional Information
Process	Displays the name and ID of the process that generated the system log message.	The information displayed in this field is different for messages generated on the local Routing Engine than for messages generated on another Routing Engine (on a system with two Routing Engines installed and operational). Messages from the other Routing Engine also include the identifiers **re0** and **re1** to identify the Routing Engine.
Severity	Severity level of a message is indicated by different colors. • **Unknown**—Gray—Indicates no severity level is specified. • **Debug/Info/Notice**—Green—Indicates conditions that are not errors but are of interest or might warrant special handling. • **Warning**—Yellow—Indicates conditions that warrant monitoring. • **Error**—Blue—Indicates standard error conditions that generally have less serious consequences than errors in the emergency, alert, and critical levels. • **Critical**—Pink—Indicates critical conditions, such as hard-drive errors. • **Alert**—Orange—Indicates conditions that require immediate correction, such as a corrupted system database. • **Emergency**—Red—Indicates system panic or other conditions that cause the switch to stop functioning.	A severity level indicates how seriously the triggering event affects switch functions. When you configure a location for logging a facility, you also specify a severity level for the facility. Only messages from the facility that are rated at that level or higher are logged to the specified file.
Event ID	Displays a code that uniquely identifies the message. The prefix on each code identifies the message source, and the rest of the code indicates the specific event or error.	The event ID begins with a prefix that indicates the generating software process. Some processes on a switch do not use codes. This field might be blank in a message generated from such a process. An event can belong to one of the following type categories: • **Error**—Indicates an error or failure condition that might require corrective action. • **Event**—Indicates a condition or occurrence that does not generally require corrective action.
Event Description	Displays a more detailed explanation of the message.	
Time	Displays the time at which the message was logged.	

Related Documentation

- Checking Active Alarms with the J-Web Interface on page 544
- Understanding Alarm Types and Severity Levels on J-EX Series Switches on page 533

Checking Active Alarms with the J-Web Interface

Purpose Use the monitoring functionality to view alarm information for the J-EX Series switches including alarm type, alarm severity, and a brief description for each active alarm on the switching platform.

Action To view the active alarms:

1. Select **Monitor> Events and Alarms > View Alarms** in the J-Web interface.
2. Select an alarm filter based on alarm type, severity, description, and date range.
3. Click **Go**.

   All the alarms matching the filter are displayed.

NOTE: When the switch is reset, the active alarms are displayed.

Meaning Table 76 on page 544 lists the alarm output fields.

Table 76: Summary of Key Alarm Output Fields

Field	Values
Type	Category of the alarm: • Chassis—Indicates an alarm condition on the chassis (typically an environmental alarm such as one related to temperature). NOTE: When you connect power to a J-EX4200 switch, the Alarm (ALM) LED lights red to indicate that the network link is disconnected. This behavior is normal. Plugging an active Ethernet cable into the management (MGMT) port on the switch completes the network link and turns off the ALM LED. Connecting the switch to a dedicated management console instead of a network does not affect the ALM LED. The LED remains red until the switch is connected to a network. • System—Indicates an alarm condition in the system.
Severity	Alarm severity—either major (red) or minor (yellow).
Description	Brief synopsis of the alarm.
Time	Date and time when the failure was detected.

Related Documentation

- Monitoring System Log Messages on page 541
- Dashboard for J-EX Series Switches on page 534
- Understanding Alarm Types and Severity Levels on J-EX Series Switches on page 533

Monitoring Chassis Alarms for a J-EX8200 Switch

Purpose This document provides information on chassis alarm conditions, and how you should respond when a certain chassis alarm is seen on your switch.

Various conditions related to the chassis components trigger yellow and red alarms. You cannot configure these conditions. See “Understanding Alarm Types and Severity Levels on J-EX Series Switches” on page 533.

Action You can monitor chassis alarms by watching the ALM chassis status LED and using the LCD panel to gather information about the alarm. See Chassis Status LEDs in a J-EX8200 Switch and LCD Panel in a J-EX8200 Switch.

To display switch chassis alarms in the CLI, use the following command

```
user@host> show chassis alarms
```

The command output displays the number of alarms currently active, the time when the alarm began, the severity level, and an alarm description. Note the date and time of an alarm so that you can correlate it with error messages in the messages system log file.

You can also monitor chassis alarms using the J-Web interface. See “Checking Active Alarms with the J-Web Interface” on page 544.

Table 77 on page 545 lists some of the chassis alarms that a J-EX8200 switch can generate.

Table 77: Chassis Alarms for J-EX8200 Switches

Component	Alarm Condition	Remedy	Severity	Additional Information
Fan tray	The fan tray has been removed from the chassis.	Install the fan tray.	Yellow/Red	The switch will eventually get too hot to operate if a fan tray is removed. Temperature alarms will follow. This alarm is expected during fan tray removal and installation.
Fan tray	One or more fans in a fan tray is spinning below the required speed.	Replace the fan tray.	Red	Individual fans cannot be replaced; you must replace the fan tray.
Fan tray	The fan tray’s internal connection to the switch is not functioning properly.	Remove and reinsert the fan tray. If removing and reinserting the fan tray does not resolve the problem, reboot the switch.	Red	The switch will eventually get too hot to operate if a fan tray is not operating. Temperature alarms will follow.

Table 77: Chassis Alarms for J-EX8200 Switches *(continued)*

Power supply	A power supply slot that contained a power supply at bootup is now empty.	Install a power supply in the empty power supply slot.	Yellow	You can ignore this alarm in cases in which a power supply slot can remain empty. You will not see this alarm if the switch is booted with an empty power supply slot. This alarm is expected during power supply removal and installation. This alarm can be triggered by a line card insertion. The alarm condition corrects itself when seen for this reason.
Power supply	A power supply has failed due to an input or output failure, or due to temperature issues.	Replace the failed power supply.	Red	
Power supply	A power supply's internal connection to the switch is not operating properly.	Remove and reinsert the power supply. If removing and reinserting the power supply does not resolve the problem, reboot the switch.	Red	
Temperature	The chassis warm temperature threshold has been exceeded and fan speeds have increased.	Adjust room temperature downward, if possible. Ensure airflow through the switch is unobstructed.	Yellow	The chassis is warm and should be cooled down. The switch is still functioning normally. To monitor temperature: user@switch> **show chassis environment** To monitor temperature thresholds: user@switch> **show chassis temperature-thresholds**

Table 77: Chassis Alarms for J-EX8200 Switches *(continued)*

Temperature	The chassis high temperature threshold has been exceeded and the fans are operating at full speed.	Adjust room temperature downward, if possible. Ensure airflow through the switch is unobstructed.	Red	The chassis is hot and should be cooled down. The switch might still function normally but is close to shutting down if it hasn't already. To monitor temperature: user@switch> **show chassis environment** To monitor temperature thresholds: user@switch> **show chassis temperature-thresholds**
Temperature	The chassis warm temperature threshold has been exceeded, and one or more fans are not operating properly. The operating fans are running at full speed.	Replace the fan tray that has the faulty fan or fans. Adjust room temperature downward, if possible. Ensure airflow through the switch is unobstructed.	Yellow	The chassis is warm and should be cooled down. The switch is still functioning normally. To monitor temperature: user@switch> **show chassis environment** To monitor temperature thresholds: user@switch> **show chassis temperature-thresholds**
Temperature	The chassis high temperature threshold has been exceeded, and one or more fans is not operating properly. The operating fans are running at full speed.	Replace the fan tray that has the faulty fan or fans. Adjust room temperature downward, if possible. Ensure airflow through the switch is unobstructed.	Red	The chassis is hot and should be cooled down. The switch might still function normally but is close to shutting down if it hasn't already. To monitor temperature: user@switch> **show chassis environment** To monitor temperature thresholds: user@switch> **show chassis temperature-thresholds**
Temperature	The temperature sensor on a hardware component has failed.	Replace the hardware component.	Yellow	
Routing Engine (RE), Switch Fabric and Routing Engine (SRE), or Switch Fabric (SF) module	The RE, SRE, or SF module has failed.	The RE, SRE, or SF module must be replaced.	Red	

Table 77: Chassis Alarms for J-EX8200 Switches *(continued)*

Link Status	The link to the network is down.	Check network connectivity.	Red or Yellow	The network link is disabled by default, so you might see this alarm before you connect the switch to the network.

Related Documentation

- Checking Active Alarms with the J-Web Interface on page 544
- Chassis Status LEDs in a J-EX8200 Switch

Monitoring Switch Control Traffic

Purpose Use the packet capture feature when you need to quickly capture and analyze switch control traffic on a switch. The packet capture feature allows you to capture traffic destined for or originating from the Routing Engine.

Action To use the packet capture feature in the J-Web interface, select **Troubleshoot>Packet Capture**.

To use the packet capture feature in the CLI, enter the following CLI command:

monitor traffic

Meaning You can use the packet capture feature to compose expressions with various matching criteria to specify the packets that you want to capture. You can decode and view the captured packets in the J-Web interface as they are captured. The packet capture feature does not capture transient traffic.

Table 78: Packet Capture Field Summary

Field	Function	Your Action
Interface	Specifies the interface on which the packets are captured. If you select default, packets on the Ethernet management port 0, are captured.	From the list, select an interface—for example, **ge-0/0/0**.
Detail level	Specifies the extent of details to be displayed for the packet headers. • Brief—Displays the minimum packet header information. This is the default. • Detail—Displays packet header information in moderate detail. • Extensive—Displays the maximum packet header information.	From the list, select **Detail**.
Packets	Specifies the number of packets to be captured. Values range from **1** to **1000**. Default is **10**. Packet capture stops capturing packets after this number is reached.	From the list, select the number of packets to be captured—for example, 10.

Table 78: Packet Capture Field Summary *(continued)*

Field	Function	Your Action
Addresses	Specifies the addresses to be matched for capturing the packets using a combination of the following parameters: • Direction—Matches the packet headers for IP address, hostname, or network address of the source, destination or both. • Type—Specifies if packet headers are matched for host address or network address. You can add multiple entries to refine the match criteria for addresses.	Select address-matching criteria. For example: 1. From the Direction list, select **source**. 2. From the Type list, select **host**. 3. In the Address box, type **10.1.40.48**. 4. Click **Add**.
Protocols	Matches the protocol for which packets are captured. You can choose to capture TCP, UDP, or ICMP packets or a combination of TCP, UDP, and ICMP packets.	From the list, select a protocol—for example, tcp.
Ports	Matches packet headers containing the specified source or destination TCP or UDP port number or port name.	Select a direction and a port. For example: • From the Type list, select **src**. • In the Port box, type **23**.
Advanced Options		
Absolute TCP Sequence	Specifies that absolute TCP sequence numbers are to be displayed for the packet headers.	To display absolute TCP sequence numbers in the packet headers, select this check box.
Layer 2 Headers	Specifies that link-layer packet headers are to be displayed.	To include link-layer packet headers while capturing packets, select this check box.
Non-Promiscuous	Specifies not to place the interface in promiscuous mode, so that the interface reads only packets addressed to it. In promiscuous mode, the interface reads every packet that reaches it.	To read all packets that reach the interface, select this check box.
Display Hex	Specifies that packet headers, except link-layer headers, are to be displayed in hexadecimal format.	To display the packet headers in hexadecimal format, select this check box.
Display ASCII and Hex	Specifies that packet headers are to be displayed in hexadecimal and ASCII format.	To display the packet headers in ASCII and hexadecimal formats, select this check box.
Header Expression	Specifies the match condition for the packets to be captured. The match conditions you specify for Addresses, Protocols, and Ports are displayed in expression format in this field.	You can enter match conditions directly in this field in expression format or modify the expression composed from the match conditions you specified for Addresses, Protocols, and Ports. If you change the match conditions specified for Addresses, Protocols, and Ports again, packet capture overwrites your changes with the new match conditions.
Packet Size	Specifies the number of bytes to be displayed for each packet. If a packet header exceeds this size, the display is truncated for the packet header. The default value is 96 bytes.	Type the number of bytes you want to capture for each packet header—for example, **256**.

Table 78: Packet Capture Field Summary *(continued)*

Field	Function	Your Action
Don't Resolve Addresses	Specifies that IP addresses are not to be resolved into hostnames in the packet headers displayed.	To prevent packet capture from resolving IP addresses to hostnames, select this check box.
No Timestamp	Suppresses the display of packet header timestamps.	To stop displaying timestamps in the captured packet headers, select this check box.
Write Packet Capture File	Writes the captured packets to a file in PCAP format in /var/tmp. The files are named with the prefix jweb-pcap and the extension .pcap. If you select this option, the decoded packet headers are not displayed on the packet capture page.	To decode and display the packet headers on the J-Web page, clear this check box.

Related Documentation

- Using the CLI Terminal on page 135

Monitoring System Properties

Purpose Use the monitoring functionality to view system properties such as the name and IP address of the switch and resource usage.

Action To monitor system properties in the J-Web interface, select **Monitor > System View > System Information**.

To monitor system properties in the CLI, enter the following commands:

- **show system uptime**
- **show system users**
- **show system storage**

Meaning Table 79 on page 550 summarizes key output fields in the system properties display.

Table 79: Summary of Key System Properties Output Fields

Field	Values	Additional Information
General Information		
Serial Number	Serial number for the switch.	
Junos OS Version	Version of Junos OS active on the switch, including whether the software is for domestic or export use.	Export software is for use outside of the U.S. and Canada.
Hostname	The name of switch.	
IP Address	The IP address of the switch.	

Table 79: Summary of Key System Properties Output Fields *(continued)*

Field	Values	Additional Information
Loopback Address	The loopback address.	
Domain Name Server	The address of the domain name server.	
Time Zone	The time zone on the switch.	
Time		
Current Time	Current system time, in Coordinated Universal Time (UTC).	
System Booted Time	Date and time when the switch was last booted and how long it has been running.	
Protocol Started Time	Date and time when the switching protocols were last started and how long they have been running.	
Last Configured Time	Date and time when a configuration was last committed. This field also shows the name of the user who issued the last **commit** command, through either the J-Web interface or the CLI.	
Load Average	The CPU load average for 1, 5, and 15 minutes.	
Storage Media		
Internal Flash Memory	Memory usage details of internal flash.	
External Flash Memory	Usage details of external flash memory.	
Logged in Users Details		
User	Username of any user logged in to the switching platform.	
Terminal	Terminal through which the user is logged in.	
From	System from which the user has logged in. A hyphen indicates that the user is logged in through the console.	
Login Time	Time when the user logged in.	This is the **LOGIN@** field in **show system users** command output.

Table 79: Summary of Key System Properties Output Fields *(continued)*

Field	Values	Additional Information
Idle Time	How long the user has been idle.	

Related Documentation

- Monitoring System Process Information on page 554
- Understanding J-Web User Interface Sessions on page 133

Monitoring Chassis Information

Purpose Use the monitoring functionality to view chassis properties such as general switch information, temperature and fan status, and resource information for the J-EX Series switch.

Action To view chassis properties in the J-Web interface, select **Monitor > System View > Chassis Information**.

To view chassis properties in the CLI, enter the following commands:

- **show chassis environment**
- **show chassis fpc**
- **show chassis hardware**

Meaning Table 80 on page 552 gives information about the key output fields for chassis information.

NOTE: For a J-EX4200 standalone switch, FPC refers to the switch itself. In a Virtual Chassis configuration, FPC refers to the member switch. In a J-EX8200 switch, FPC refers to the line card.

Table 80: Summary of the Key Output Fields for Chassis Information

Field	Values
Routing Engine Details	Select the **Master** tab to view details about the master Routing Engine or select **Backup** to view details about the backup Routing Engine.

Table 80: Summary of the Key Output Fields for Chassis Information *(continued)*

Field	Values
Name/Value	This table displays the following details of the master Routing Engine: • Routing engine module • Model • Version • Part number • Serial number • Memory utilization • Temperature • Start time • CPU load average for 1, 5, and 15 minutes
Power and Fan Tray Details	
Power	Select the **Power** tab to view details of the power supplies.
Name/Value	Displays the status and model number of each power supply.
Fan	Select the **Fan** tab to view details about the fans.
Name/Value	Displays the status of each fan in the corresponding FPC.
Chassis Component Details	
Select component	Select an FPC to view **General**, **Temperature**, **Resource**, and **Sub-component** details.
General	Select the **General** tab to view the general information about the chassis components.
Name/Value	Displays general information: • Version—Revision level. Supply the version number when reporting hardware problems to customer support. • Part Number • Serial Number—Supply the serial number when contacting customer support about the switch chassis. • Description—Brief text description.
Temperature	Select the **Temperature** tab to view the temperature details of the components in the selected FPC.
Name/Value	Displays the temperature details of the sensors present in the selected FPC.
Resource	Select the **Resource** tab to view the resource details of the selected FPC.

Table 80: Summary of the Key Output Fields for Chassis Information *(continued)*

Field	Values

Name/Value	Displays resource details:
	• State:
	  • **Dead**—Held in reset because of errors.
	  • **Diag**—The FPC is running diagnostics.
	  • **Dormant**—Held in reset.
	  • **Empty**—No FPC is present.
	  • **Online**—The FPC is online and running.
	  • **Probed**—Probe is complete. The FPC is awaiting restart of the Packet Forwarding Engine (PFE).
	  • **Probe-wait**—The FPC is waiting for the probe operation to start.
	• **Total CPU DRAM**—Total DRAM, in megabytes, available to the FPC.
	• **Start time**—Date and time the switch was last rebooted.

Related Documentation

- Monitoring System Process Information on page 554
- Monitoring System Properties on page 550
- Dashboard for J-EX Series Switches on page 534

Monitoring System Process Information

Purpose Use the monitoring functionality to view the processes running on the switch.

Action To view the software processes running on the switch in the J-Web interface, select **Monitor>System View>Process Details**.

To view the software processes running on the switch in the CLI, enter the following command.

show system processes

Meaning Table 81 on page 555 summarizes the output fields in the system process information display.

The display includes the total CPU load and total memory utilization.

Table 81: Summary of System Process Information Output Fields

Field	Values
PID	Identifier of the process.
Name	Owner of the process.
State	Current state of the process.
CPU Load	Percentage of the CPU that is being used by the process.
Memory Utilization	Amount of memory that is being used by the process.
Start Time	Time of day when the process started.

Related Documentation

- Monitoring System Properties on page 550
- For more information about show system properties command, see **show system uptime on page 842**

Managing Log, Temporary, and Crash Files on the Switch (J-Web Procedure)

You can use the J-Web interface to rotate log files and delete unnecessary log, temporary, and crash files on the switching platform.

1. Cleaning Up Files on page 555
2. Downloading Files on page 556
3. Deleting Files on page 556

Cleaning Up Files

If you are running low on storage space, use the file cleanup procedure to quickly identify files to delete.

The file cleanup procedure performs the following tasks:

- Rotates log files—Archives the current log files, and creates fresh log files.
- Deletes log files in /var/log—Deletes files that are not currently being written to.
- Deletes temporary files in /var/tmp—Deletes files that have not been accessed within two days.
- Deletes all crash files in /var/crash—Deletes core files that the switch has written during an error.

To rotate log files and delete unnecessary files with the J-Web interface:

1. Select **Maintain>Files**.

2. In the Clean Up Files section, click **Clean Up Files**. The switching platform rotates log files and identifies files that can be safely deleted.

   The J-Web interface displays the files that you can delete and the amount of space that will be freed on the file system.

3. Click one:

   - To delete the files and return to the Files page, click **OK**.
   - To cancel your entries and return to the list of files in the directory, click **Cancel**.

Downloading Files

You can use the J-Web interface to download a copy of an individual log, temporary, or crash file from the switching platform. When you download a file, it is not deleted from the file system.

To download files with the J-Web interface:

1. In the J-Web interface, select **Maintain>Files**.
2. In the Download and Delete Files section, click one:

   - Log Files—Log files in the /var/log directory on the switch.
   - Temporary Files—Lists the temporary files in the /var/tmp directory on the switching platform.
   - Jailed Temporary Files (Install, Session, etc)—Lists the files in the /var/jail/tmp directory on the switching platform.
   - Crash (Core) Files—Lists the core files in the /var/crash directory on the switching platform.

     The J-Web interface displays the files located in the directory.

3. Select the files that you want to download and click **Download**.
4. Choose a location for the saved file.

   The file is saved as a text file, with a **.txt** file extension.

Deleting Files

You can use the J-Web interface to delete an individual log, temporary, and crash file from the switching platform. When you delete the file, it is permanently removed from the file system.

CAUTION: If you are unsure whether to delete a file from the switching platform, we recommend using the Clean Up Files tool described in Cleaning Up Files. This tool determines which files can be safely deleted from the file system.

To delete files with the J-Web interface:

1. Select **Maintain>Files**.
2. In the Download and Delete Files section, click one:
   - Log Files—Lists the log files in the /var/log directory on the switching platform.
   - Temporary Files—Lists the temporary files in the /var/tmp directory on the switching platform.
   - Jailed Temporary Files (Install, Session, etc)—Lists the files in the /var/jail/tmp directory on the switching platform.
   - Crash (Core) Files—Lists the core files in the /var/crash directory on the switching platform.

     The J-Web interface displays the files in the directory.
3. Select the box next to each file you plan to delete.
4. Click **Delete**.

   The J-Web interface displays the files you can delete and the amount of space that will be freed on the file system.
5. Click one of the following buttons on the confirmation page:
   - To delete the files and return to the Files page, click **OK**.
   - To cancel your entries and return to the list of files in the directory, click **Cancel**.

CHAPTER 41

Configuration Statements for System Monitoring

archive (All System Log Files)

Syntax archive <files *number*> <size *size* <world-readable | no-world-readable>;

Hierarchy Level [edit system **syslog**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure archiving properties for all system log files.

Options

files ***number***—Maximum number of archived log files to retain. When the Junos OS logging utility has written a defined maximum amount of data to a log file ***logfile***, it closes the file, compresses it, and renames it to ***logfile*.0.gz** (the amount of data is determined by the **size** statement at this hierarchy level). The utility then opens and writes to a new file called ***logfile***. When the new file reaches the maximum size, the ***logfile*.0.gz** file is renamed to ***logfile*.1.gz**, and the new file is closed, compressed, and renamed ***logfile*.0.gz**. By default, the logging facility creates up to ten archive files in this manner. Once the maximum number of archive files exists, each time the active log file reaches the maximum size, the contents of the oldest archive file are lost (overwritten by the next oldest file).

Range: 1 through 1000

Default: 10 files

size ***size***—Maximum amount of data that the Junos OS logging utility writes to a log file ***logfile*** before archiving it (closing it, compressing it, and changing its name to ***logfile*.0.gz**). The utility then opens and writes to a new file called ***logfile***.

Syntax: ***x*k** to specify the number of kilobytes, ***x*m** for the number of megabytes, or ***x*g** for the number of gigabytes

Range: 64 KB through 1 GB

world-readable | no-world-readable—Grant all users permission to read archived log files, or restrict the permission only to the **root** user and users who have the Junos OS **maintenance** permission.

Default: **no-world-readable**

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Specifying Log File Size, Number, and Archiving Properties

archive-sites

Syntax

```
archive-sites {
  url <password password>;
}
```

Hierarchy Level [edit **event-options destinations** *destination-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify an archive site to which files are transferred. If you specify more than one archive site, the router or switch attempts to transfer to the first archive site in the list, moving to the next site only if the transfer fails.

Options

url—The archive destination specified as Hypertext Transfer Protocol (HTTP) URL, FTP URL, or secure copy (scp)-style remote file specification. URLs of the type **file://** are not supported; however, local router or switch directories are supported (for example, **/var/tmp/**).

password *password*—A plain-text password for login into the archive site.

Required Privilege Level maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation

- Defining Destinations for File Archiving by Event Policies

arguments

Syntax

```
arguments {
  argument-name argument-value;
}
```

Hierarchy Level [edit **event-options policy** *policy-name* **then event-script** *filename*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define command-line arguments for an event script that is invoked from an event policy.

Options

argument-name—Name of the argument.

argument-value—Value of the argument.

Required Privilege Level maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation

- Executing Event Scripts in an Event Policy

attributes-match

Syntax

```
attributes-match {
  event1.attribute-name equals event2.attribute-name;
  event.attribute-name matches regular-expression;
  event1.attribute-name starts-with event2.attribute-name;
}
```

Hierarchy Level [edit **event-options policy** *policy-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Execute the policy only if the attributes of two events are correlated or if the attribute of one event matches a regular expression.

If the **attributes-match** statement includes the **equals** or **starts-with** options, or if it includes a **matches** option that includes a clause for an event that is not specified at the **[edit event-options policy *policy-name* events]** hierarchy level, you must include one or more **within** statements in the same policy configuration.

The statements are explained separately.

Required Privilege Level maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation

- Using Correlated Events to Trigger an Event Policy
- **within on page 608**

commands

Syntax

```
commands {
  "command";
}
```

Hierarchy Level [edit **event-options policy** *policy-name* **then execute-commands**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify an operational mode command to be issued on receipt of an event.

Options ***command***—Command to be issued. Enclose each command in quotation marks (“ ”). The event process (eventd) issues the commands in the order in which they appear in the configuration.

You can include variables in commands. The eventd process replaces each variable with values contained in the event that triggers the policy. You can use command variables of the following forms:

- **{$$.*attribute-name*}**—The double dollar sign (**$$**) notation represents the event that is triggering a policy. When combined with an attribute name, the command variable is replaced by the value of the attribute name of the triggering event.
- **{$*event*.*attribute-name*}**—The dollar sign with the event name (**$*event***) notation represents the most recent event that matches the specified event. The variable is replaced by the value of the attribute name of the most recent event that matches ***event***.
- **{$*.*attribute-name*}**—The dollar sign with the asterisk (**$***) notation represents the most recent event that matches any of the correlating events. The variable is replaced by the value of the attribute name of the most recent event that matches any of the events specified in the policy configuration.

Required Privilege Level maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation

- Configuring an Event Policy to Execute Operational Mode Commands
- Representing the Correlating Event in an Event Policy

console (System Logging)

Syntax

```
console {
  facility severity;
}
```

Hierarchy Level [edit system syslog]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the logging of system messages to the system console.

Options

facility—Class of messages to log. To specify multiple classes, include multiple ***facility severity*** statements. For a list of the facilities, see Junos OS System Logging Facilities and Message Severity Levels.

severity—Severity of the messages that belong to the facility specified by the paired ***facility*** name. Messages with severities the specified level and higher are logged. For a list of the severities, see Junos OS System Logging Facilities and Message Severity Levels.

Required Privilege Level
system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Directing System Log Messages to the Console
- *Junos OS System Log Messages Reference*

destination

Syntax

```
destination destination-name {
    retry-count count retry-interval seconds;
    transfer-delay seconds;
}
```

Hierarchy Level [edit **event-options policy** *policy-name* **then event-script** *filename*],
[edit **event-options policy** *policy-name* **then execute-commands**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Assign a location to which to upload command or script output for the specified policy.

Options ***destination-name***—Name of a destination defined in the **destinations** statement at the **[edit event-options]** hierarchy level.

The remaining statements are explained separately.

Required Privilege Level maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation

- Configuring an Event Policy to Execute Operational Mode Commands
- Executing Event Scripts in an Event Policy
- **destinations on page 566**

destinations

Syntax

```
destinations {
  destination-name {
    archive-sites {
      url <password password>;
    }
    transfer-delay seconds;
  }
}
```

Hierarchy Level [edit event-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define one or more destinations, each with a unique name and other attributes. You can use the destination as a storage location for command output and for various files, such as system log files and core files.

Options ***destination-name***—Name of a destination.

The remaining statements are explained separately.

Required Privilege Level maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation
- Defining Destinations for File Archiving by Event Policies

equals

Syntax *event1.attribute-name* equals *event2.attribute-name*;

Hierarchy Level [edit event-options policy *policy-name* attributes-match]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Execute the policy only if the specified attribute of ***event1*** equals the specified attribute of ***event2***.

Options ***event1.attribute-name***—Attribute of one event.

event2.attribute-name—Attribute of another event.

Required Privilege Level maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation
- Using Correlated Events to Trigger an Event Policy

event-options

Syntax

```
event-options {
  destinations {
    destination-name {
      archive-sites {
        url <password password>;
      }
      transfer-delay seconds;
    }
  }
  event-script {
    file filename {
      checksum (md5 | sha-256 | sha1) hash;
      refresh;
      refresh-from url;
      remote-execution {
        remote-hostname {
          passphrase user-password;
          username user-login;
        }
      }
      source url;
    }
    refresh;
    refresh-from url;
    traceoptions {
      file <filename> <files number> <size size> <world-readable | no-world-readable>;
      flag flag;
      no-remote-trace;
    }
  }
  generate-event event-name {
    time-interval seconds;
    time-of-day hh:mm:ss;
  }
  policy policy-name {
    attributes-match {
      event1.attribute-name equals event2.attribute-name;
      event.attribute-name matches regular-expression;
      event1.attribute-name starts-with event2.attribute-name;
    }
    events [ events ];
    within seconds not events [ events ];
    then {
      event-script filename {
        arguments {
          argument-name argument-value;
        }
        output-filename filename;
        destination destination-name {
          retry-count count retry-interval seconds;
          transfer-delay seconds;
        }
```

```
        }
        execute-commands {
          commands {
            "command";
          }
          destination destination-name {
            retry-count count retry-interval seconds;
            transfer-delay seconds;
          }
          output-filename filename;
          output-format (text | xml);
          user-name username;
        }
        ignore;
        raise-trap;
        upload filename (filename | committed) destination destination-name {
          retry-count count retry-interval seconds;
          transfer-delay seconds;
          user-name username;
        }
      }
    }
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>;
      flag flag;
    }
}
```

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure event policies.

The statements are explained separately.

Required Privilege Level maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

events (Associating Events with a Policy)

Syntax events [*events*];

Hierarchy Level [edit **event-options policy** *policy-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Create a list of events that trigger this policy. If one or more of the listed events occurs, the policy is executed.

Options **[*events*]**—List of events. Events can be internally generated, or they can be generated by Junos OS processes.

Required Privilege Level maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation
- Using Correlated Events to Trigger an Event Policy

events (Correlating Events with Each Other)

Syntax events [*events*];

Hierarchy Level [edit **event-options policy** *policy-name* **within** *seconds*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Create a list of events that must occur within a specified time interval for the policy to be triggered.

Options **[*events*]**—List of events. Events can be internally generated, or they can be generated by Junos OS processes.

Required Privilege Level maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation
- Using Correlated Events to Trigger an Event Policy

event-script

Syntax

```
event-script {
  file filename {
    checksum (md5 | sha-256 | sha1) hash;
    refresh;
    refresh-from url;
    remote-execution {
      remote-hostname {
        passphrase user-password;
        username user-login;
      }
    }
    source url;
  }
  refresh;
  refresh-from url;
  traceoptions {
    file <filename> <files number> <size size> <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
  }
}
```

Hierarchy Level [edit event-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For Junos OS event scripts, configure scripting mechanisms.

The statements are explained separately.

Required Privilege Level maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation

- Implementing Event Scripts

event-script

Syntax

```
event-script filename {
  arguments {
    argument-name argument-value;
  }
  destination destination-name {
    retry-count count retry-interval seconds;
    transfer-delay seconds;
  }
  output-filename filename;
  output-format (text | xml);
  user-name username;
}
```

Hierarchy Level [edit **event-options policy** *policy-name* **then**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description On receipt of an event, specify operational mode commands to be issued, the format of the command output, and a name and destination for the output file.

The statements are explained separately.

Required Privilege Level maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation

- Executing Event Scripts in an Event Policy

execute-commands

Syntax

```
execute-commands {
  commands {
    "command";
  }
  destination destination-name {
    retry-count count retry-interval seconds;
    transfer-delay seconds;
  }
  output-filename filename;
  output-format (text | xml);
  user-name username;
}
```

Hierarchy Level [edit **event-options policy** *policy-name* **then**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description On receipt of an event, specify operational mode commands to be issued, the format of the command output, and a name and destination for the output file.

The statements are explained separately.

Required Privilege Level maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation

- Configuring an Event Policy to Execute Operational Mode Commands

explicit-priority

Syntax

```
explicit-priority;
```

Hierarchy Level [edit system syslog **file** *filename*],
[edit system syslog **host**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Record the priority (facility and severity level) in each standard-format system log message directed to a file or remote destination.

When the **structured-data** statement is also included at the **[edit system syslog file *filename*]** hierarchy level, this statement is ignored for the file.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Including Priority Information in System Log Messages
- *Junos OS System Log Messages Reference*
- **structured-data on page 590**

facility-override

Syntax facility-override *facility*;

Hierarchy Level [edit system syslog **host**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Substitute an alternate facility for the default facilities used when messages are directed to a remote destination.

Options ***facility***—Alternate facility to substitute for the default facilities. For a list of the possible facilities, see Junos OS System Log Alternate Facilities for Remote Logging.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Changing the Alternative Facility Name for Remote System Log Messages
- *Junos OS System Log Messages Reference*

file

Syntax

```
file filename {
  checksum (md5 | sha-256 | sha1) hash;
  refresh;
  refresh-from url;
  remote-execution {
    remote-hostname {
      passphrase user-password;
      username user-login;
    }
  }
  source url;
}
```

Hierarchy Level [edit event-options event-script]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For Junos OS event scripts, enable an event script that is located in the **/var/db/scripts/event** directory.

Options ***filename***—The name of an Extensible Stylesheet Language Transformations (XSLT) or Stylesheet Language Alternative Syntax (SLAX) file containing an event script.

The statements are explained separately.

Required Privilege Level maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation

- Enabling an Event Script

file (System Logging)

Syntax

```
file filename {
  facility severity;
  archive {
    files number;
    size size;
    (no-world-readable | world-readable);
  }
  explicit-priority;
  match "regular-expression";
  structured-data {
    brief;
  }
}
```

Hierarchy Level [edit system syslog]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the logging of system messages to a file.

Options

facility—Class of messages to log. To specify multiple classes, include multiple ***facility severity*** statements. For a list of the facilities, see Junos OS System Logging Facilities and Message Severity Levels.

file *filename*—File in the /var/log directory in which to log messages from the specified facility. To log messages to more than one file, include more than one **file** statement.

severity—Severity of the messages that belong to the facility specified by the paired ***facility*** name. Messages with severities the specified level and higher are logged. For a list of the severities, see Junos OS System Logging Facilities and Message Severity Levels.

The remaining statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Directing System Log Messages to a Log File
- *Junos OS System Log Messages Reference*

files

Syntax files *number*;

Hierarchy Level [edit system syslog **archive**],
[edit system syslog file *filename* archive]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the maximum number of archived log files to retain. When the Junos OS logging utility has written a defined maximum amount of data to a log file ***logfile***, it closes the file, compresses it, and renames it to ***logfile*.0.gz** (for information about the maximum file size, see **size**). The utility then opens and writes to a new file called ***logfile***. When the new file reaches the maximum size, the ***logfile*.0.gz** file is renamed to ***logfile*.1.gz**, and the new file is closed, compressed, and renamed ***logfile*.0.gz**. By default, the logging facility creates up to ten archive files in this manner. Once the maximum number of archive files exists, each time the active log file reaches the maximum size, the contents of the oldest archive file are lost (overwritten by the next oldest file).

Options ***number***—Maximum number of archived files.

Range: 1 through 1000

Default: 10 files

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Specifying Log File Size, Number, and Archiving Properties
- *Junos OS System Log Messages Reference*
- **size on page 588**

generate-event

Syntax

```
generate-event event-name {
    time-interval seconds;
    time-of-day hh:mm:ss;
}
```

Hierarchy Level [edit event-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Generate an internal event, based on a time interval or the time of day.

Options ***event-name***—Name of an internally generated event.

The statements are explained separately.

Required Privilege Level maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation

- Generating Internal Events to Trigger Event Policies

host

Syntax

```
host (hostname | other-routing-engine) {
  facility severity;
  explicit-priority;
  facility-override facility;
  log-prefix string;
  match "regular-expression";
}
```

J-EX Series Switches

```
host (hostname | other-routing-engine | scc-master) {
  facility severity;
  explicit-priority;
  facility-override facility;
  log-prefix string;
  match "regular-expression";
}
```

TX Matrix Plus Router

```
host (hostname | other-routing-engine | sfc0-master) {
  facility severity;
  explicit-priority;
  facility-override facility;
  log-prefix string;
  match "regular-expression";
}
```

Hierarchy Level [edit system syslog]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the logging of system messages to a remote destination.

Options

facility—Class of messages to log. To specify multiple classes, include multiple ***facility severity*** statements. For a list of the facilities, see Junos OS System Logging Facilities and Message Severity Levels.

hostname—IPv4 address, IPv6 address, or fully qualified hostname of the remote machine to which to direct messages. To direct messages to multiple remote machines, include a **host** statement for each one.

other-routing-engine—Direct messages to the other Routing Engine on a router or switch with two Routing Engines installed and operational.

severity—Severity of the messages that belong to the facility specified by the paired ***facility*** name. Messages with severities the specified level and higher are logged. For a list of the severities, see Junos OS System Logging Facilities and Message Severity Levels.

The remaining statements are explained separately.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Directing System Log Messages to a Remote Machine or the Other Routing Engine
- *Junos OS System Log Messages Reference*

ignore

Syntax
```
ignore;
```

Hierarchy Level
[edit **event-options policy** *policy-name* **then**]

Release Information
Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Define a policy that ignores particular events. If one or more of the listed events occur, a system log message for the event is not generated, and no further policies associated with this event are processed. If you include the **ignore** statement in a policy configuration, you cannot configure any other actions in the policy.

Required Privilege Level
maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation
- Configuring Event Policies to Ignore an Event

interface (Accounting or Sampling)

Syntax
```
interface interface-name {
    engine-id number;
    engine-type number;
    source-address address;
}
```

Hierarchy Level
[edit forwarding-options accounting *name* output],
[edit forwarding-optionssamplingoutput]

Release Information
Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Specify the output interface for monitored traffic.

Options
interface-name—Name of the interface.

The remaining statements are explained separately.

Usage Guidelines
See Configuring Discard Accounting or Configuring Traffic Sampling.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

log-prefix

Syntax log-prefix *string*;

Hierarchy Level [edit system syslog **host**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Include a text string in each message directed to a remote destination.

Options ***string***—Text string to include in each message.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Adding a Text String to System Log Messages
- *Junos OS System Log Messages Reference*

match

Syntax match "*regular-expression*";

Hierarchy Level [edit system syslog **file** *filename*],
[edit system syslog **host** *hostname* | other-routing-engine| scc-master)],
[edit system syslog **user** (*username* | *)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify a text string that must (or must not) appear in a message for the message to be logged to a destination.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Using Regular Expressions to Refine the Set of Logged Messages

not

Syntax not events [*events*];

Hierarchy Level [edit **event-options policy** *policy-name* **within** *seconds*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Create a list of events that must not occur within the specified time interval for the policy to be triggered.

Options **[*events*]**—List of events. Events can be internally generated, or they can be generated by Junos OS processes.

Required Privilege Level maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation
- Using Correlated Events to Trigger an Event Policy

output-filename

Syntax output-filename *filename*;

Hierarchy Level [edit **event-options policy** *policy-name* **then event-script** *filename*],
[edit **event-options policy** *policy-name* **then execute-commands**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Assign a filename to which to write command or script output for the specified commands or script. For op scripts, this statement is optional.

Options ***filename***—Name of a file in which to write command or script output.

Required Privilege Level maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation
- Configuring an Event Policy to Execute Operational Mode Commands
- Executing Event Scripts in an Event Policy

output-format

Syntax output-format (text | xml);

Hierarchy Level [edit **event-options policy** *policy-name* **then event-script** *filename*],
[edit **event-options policy** *policy-name* **then execute-commands**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the format (ASCII text or XML) for the output of the specified commands or script.

Options **text**—Formatted ASCII text.

xml—Junos Extensible Markup Language (XML) tags.

Default: **xml** at the **[edit event-options policy *policy-name* then execute-commands]** hierarchy level and **text** at the **[edit event-options policy *policy-name* then event-script *filename*]** hierarchy level.

Required Privilege Level maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation

- Configuring an Event Policy to Execute Operational Mode Commands
- Executing Event Scripts in an Event Policy

policy

Syntax

```
policy policy-name {
  attributes-match {
    event1.attribute-name equals event2.attribute-name;
    event.attribute-name matches regular-expression;
    event1.attribute-name starts-with event2.attribute-name;
  }
  events [ events ];
  then {
    ... the then subhierarchy appears at the end of the [edit event-options policy policy-name]
      hierarchy level ...
  }
  within seconds {
    events [ events ];
    not events [ events ];
    trigger (on | after | until) event-count;
  }

  then {
    event-script filename {
      arguments {
        argument-name argument-value;
      }
      destination destination-name {
        retry-count count retry-interval seconds;
        transfer-delay seconds;
      }
      output-filename filename;
      output-format (text | xml);
      user-name username;
    }
    execute-commands {
      commands {
        "command";
      }
      destination destination-name {
        retry-count count retry-interval seconds;
        transfer-delay seconds;
      }
      output-filename filename;
      output-format (text | xml);
      user-name username;
    }
    ignore;
    raise-trap;
    upload filename (filename | committed) destination destination-name {
      retry-count count retry-interval seconds;
      transfer-delay seconds;
      user-name username;
    }
  }
}
```

Hierarchy Level [edit **event-options**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define an event policy to be processed by the eventd process. If you configure a policy, the **events** and **then** statements are mandatory.

You can configure multiple policies to be processed for an event. The policies are executed in the order in which they appear in the configuration. If you configure more than one policy for an event, and if one of the policies is to ignore the event, no policies that follow the **ignore** statement are executed.

Default If you do not configure a policy for an event, the event is recorded in the system log.

Options ***policy-name***—Name of an event policy.

The statements are explained separately.

Required Privilege Level maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

raise-trap

Syntax raise-trap;

Hierarchy Level [edit **event-options policy** *policy-name* **then**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define a policy that raises an SNMP trap in response to an event. If one or more of the listed events occur, the system log message for the event is converted into a trap. This enables an agent to notify a trap-based network management system (NMS) of significant events.

Required Privilege Level maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation

- Configuring Event Policies to Raise SNMP Traps

refresh

Syntax refresh;

Hierarchy Level [edit **event-options event-script**],
[edit **event-options event-script file** *filename*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For Junos OS event scripts, overwrite the local copy of all enabled event scripts or a single enabled script located in the **/var/db/scripts/event** directory with the copy located at the source URL, specified in the **source** statement at the same hierarchy level.

Required Privilege Level maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation
- Updating an Event Script from the Master Source
- **refresh-from on page 585**
- **source on page 589**

refresh-from

Syntax refresh-from *url*;

Hierarchy Level [edit **event-options event-script**],
[edit **event-options event-script file** *filename*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For Junos OS event scripts, overwrite the local copy of all enabled event scripts or a single enabled script located in the **/var/db/scripts/event** directory with the copy located at a URL other than the URL specified in the **source** statement.

Options *url*—Source specified as a Hypertext Transfer Protocol (HTTP) URL, FTP URL, or secure copy (scp)-style remote file specification.

Required Privilege Level maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation
- Updating an Event Script from an Alternate Location
- **refresh on page 585**
- **source on page 589**

remote-execution

Syntax

```
remote-execution {
  remote-hostname {
    passphrase user-password;
    username user-login;
  }
}
```

Hierarchy Level [edit **event-options event-script file** *filename*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For Junos OS event scripts, enable event scripts to invoke RPCs on a local or remote host.

Options

passphrase ***user-password***—User's password for the remote host.

remote-hostname—Name of the remote host with which the event script will communicate.

username ***username***—User's login name for the remote host.

Required Privilege Level

maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation

- Using RPCs and Operational Mode Commands in Event Scripts

retry-count

Syntax retry-count *number* retry-interval *seconds*;

Hierarchy Level [edit **event-options policy** *policy-name* **then event-script** *filename* **destination** *destination-name*],
[edit **event-options policy** *policy-name* **then execute-commands destination** *destination-name*],
[edit **event-options policy** *policy-name* **then upload** filename (*filename* | committed) destination *destination-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure an event policy to retry a file upload operation if the first attempt fails.

Default If you do not include this statement, the file upload operation is attempted one time only.

Options ***number***—Number of retries.

retry-interval *seconds*—Length of time to wait between retries.

Required Privilege Level maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation

- Configuring an Event Policy to Retry the File Upload Action

size

Syntax size *size*;

Hierarchy Level [edit system syslog **archive**],
[edit system syslog file *filename* archive]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the maximum amount of data that the Junos OS logging utility writes to a log file ***logfile*** before archiving it (closing it, compressing it, and changing its name to ***logfile*.0.gz**). The utility then opens and writes to a new file called ***logfile***. For information about the number of archive files that the utility creates in this way, see **files**.

Options ***size***—Maximum size of each system log file, in kilobytes (KB), megabytes (MB), or gigabytes (GB).

Syntax: ***x*k** to specify the number of kilobytes, ***x*m** for the number of megabytes, or ***x*g** for the number of gigabytes

Range: 64 KB through 1 GB

Default: 1 MB for MX Series routers

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Specifying Log File Size, Number, and Archiving Properties
- *Junos OS System Log Messages Reference*
- **files on page 576**

source

Syntax source *url*;

Hierarchy Level [edit **event-options event-script file** *filename*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For Junos OS event scripts, specify the location of the source file for an enabled script located in the **/var/db/scripts/event** directory. When you include the **refresh** statement at the same hierarchy level, the local copy is overwritten by the version stored at the specified URL.

Options *url*—Master source file for an event script specified as an HTTP URL, FTP URL, or scp-style remote file specification.

Required Privilege Level maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation

- **refresh on page 585**
- **refresh-from on page 585**
- Specifying a Master Source for an Event Script

structured-data

Syntax

```
structured-data {
  brief;
}
```

Hierarchy Level [edit system syslog file *filename*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Write system log messages to the log file in structured-data format, which complies with Internet draft draft-ietf-syslog-protocol-23, *The syslog Protocol* (http://tools.ietf.org/html/draft-ietf-syslog-protocol-23).

NOTE: When this statement is included, other statements that specify the format for messages written to the file are ignored (the **explicit-priority** statement at the **[edit system syslog file *filename*]** hierarchy level and the **time-format** statement at the **[edit system syslog]** hierarchy level).

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Logging Messages in Structured-Data Format
- *Junos OS System Log Messages Reference*
- **explicit-priority on page 572**
- **time-format on page 594**

syslog

Syntax

```
syslog {
  archive {
    files number;
    size maximum-file-size;
    start-time "YYYY-MM-DD.hh:mm";
    transfer-interval minutes;
    (world-readable | no-world-readable);
  }
  console {
    facility severity;
  }
  file filename {
    facility severity;
    explicit-priority;
    match "regular-expression";
    archive {
      files number;
      size maximum-file-size;
      start-time "YYYY-MM-DD.hh:mm";
      transfer-interval minutes;
      (world-readable | no-world-readable);
    }
    structured-data {
      brief;
    }
  }
  host (hostname | other-routing-engine | scc-master) {
    facility severity;
    explicit-priority;
    facility-override facility;
    log-prefix string;
    match "regular-expression";
  }
  source-address source-address;
  time-format (millisecond | year | year millisecond);
  user (username | *) {
    facility severity;
    match "regular-expression";
  }
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the types of system log messages to log to files, a remote destination, user terminals, or the system console.

The remaining statements are explained separately.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Junos OS System Log Configuration Overview
- *Junos OS System Log Messages Reference*

then

Syntax

```
then {
  event-script filename {
    arguments {
      argument-name argument-value;
    }
    destination destination-name {
      retry-count count retry-interval seconds;
      transfer-delay seconds;
    }
    output-filename filename;
    output-format (text | xml);
    user-name username;
  }
  execute-commands {
    commands {
      "command";
    }
    destination destination-name {
      retry-count count retry-interval seconds;
      transfer-delay seconds;
    }
    output-filename filename;
    output-format (text | xml);
    user-name username;
  }
  ignore;
  raise-trap;
  upload filename (filename | committed) destination destination-name {
    retry-count count retry-interval seconds;
    transfer-delay seconds;
    user-name username;
  }
}
```

Hierarchy Level [edit event-options policy *policy-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define actions to take if an event occurs. For each policy, you can configure multiple actions.

The statements are explained separately.

Required Privilege Level maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation

- Configuring an Event Policy to Upload Files
- Configuring an Event Policy to Execute Operational Mode Commands
- Executing Event Scripts in an Event Policy

- Configuring Event Policies to Ignore an Event
- Configuring Event Policies to Raise SNMP Traps

time-format

Syntax time-format (year | millisecond | year millisecond);

Hierarchy Level [edit system syslog]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Include the year, the millisecond, or both, in the timestamp on every standard-format system log message. The additional information is included for messages directed to each destination configured by a **file**, **console**, or **user** statement at the **[edit system syslog]** hierarchy level, but not to destinations configured by a **host** statement.

By default, the timestamp specifies the month, date, hour, minute, and second when the message was logged—for example, **Aug 21 12:36:30**.

NOTE: When the **structured-data** statement is included at the **[edit system syslog file *filename*]** hierarchy level, this statement is ignored for the file.

Options **millisecond**—Include the millisecond in the timestamp.

year—Include the year in the timestamp.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Including the Year or Millisecond in Timestamps
- *Junos OS System Log Messages Reference*
- **structured-data on page 590**

time-interval

Syntax time-interval *seconds*;

Hierarchy Level [edit **event-options generate-event** *event-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a frequency at which to generate a particular event.

Options ***seconds***—Time interval between internally generated events.
Range: 60 through 604,800 seconds

Required Privilege Level maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation
- Generating Internal Events to Trigger Event Policies

time-of-day

Syntax time-of-day *hh*:*mm*:*ss*;

Hierarchy Level [edit **event-options generate-event** *event-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a time of day at which to generate a particular event.

Options ***hh*:*mm*:*ss***—Time of day at which to generate an event.

Required Privilege Level maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation
- Generating Internal Events to Trigger Event Policies

time-zone

Syntax time-zone (GMT *hour-offset* | *time-zone*);

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set the local time zone. To have the time zone change take effect for all processes running on the router or switch, you must reboot the router or switch.

Default UTC

Options GMT *hour-offset*—Set the time zone relative to UTC time.

Range: –14 through +12

Default: 0

time-zone—Specify the time zone as **UTC**, which is the default time zone, or as a string such as PDT (Pacific Daylight Time), or use one of the following continents and major cities:

Africa/Abidjan, Africa/Accra, Africa/Addis_Ababa, Africa/Algiers, Africa/Asmera, Africa/Bamako, Africa/Bangui, Africa/Banjul, Africa/Bissau, Africa/Blantyre, Africa/Brazzaville, Africa/Bujumbura, Africa/Cairo, Africa/Casablanca, Africa/Ceuta, Africa/Conakry, Africa/Dakar, Africa/Dar_es_Salaam, Africa/Djibouti, Africa/Douala, Africa/El_Aaiun, Africa/Freetown, Africa/Gaborone, Africa/Harare, Africa/Johannesburg, Africa/Kampala, Africa/Khartoum, Africa/Kigali, Africa/Kinshasa, Africa/Lagos, Africa/Libreville, Africa/Lome, Africa/Luanda, Africa/Lubumbashi, Africa/Lusaka, Africa/Malabo, Africa/Maputo, Africa/Maseru, Africa/Mbabane, Africa/Mogadishu, Africa/Monrovia, Africa/Nairobi, Africa/Ndjamena, Africa/Niamey, Africa/Nouakchott, Africa/Ouagadougou, Africa/Porto-Novo, Africa/Sao_Tome, Africa/Timbuktu, Africa/Tripoli, Africa/Tunis, Africa/Windhoek

America/Adak, America/Anchorage, America/Anguilla, America/Antigua, America/Aruba, America/Asuncion, America/Barbados, America/Belize, America/Bogota, America/Boise, America/Buenos_Aires, America/Caracas, America/Catamarca, America/Cayenne, America/Cayman, America/Chicago, America/Cordoba, America/Costa_Rica, America/Cuiaba, America/Curacao, America/Dawson, America/Dawson_Creek, America/Denver, America/Detroit, America/Dominica, America/Edmonton, America/El_Salvador, America/Ensenada, America/Fortaleza, America/Glace_Bay, America/Godthab, America/Goose_Bay, America/Grand_Turk, America/Grenada, America/Guadeloupe, America/Guatemala, America/Guayaquil, America/Guyana, America/Halifax, America/Havana, America/Indiana/Knox, America/Indiana/Marengo, America/Indiana/Vevay, America/Indianapolis, America/Inuvik, America/Iqaluit, America/Jamaica, America/Jujuy, America/Juneau, America/La_Paz, America/Lima, America/Los_Angeles, America/Louisville, America/Maceio, America/Managua, America/Manaus, America/Martinique, America/Mazatlan, America/Mendoza, America/Menominee, America/Mexico_City, America/Miquelon, America/Montevideo, America/Montreal, America/Montserrat, America/Nassau, America/New_York, America/Nipigon, America/Nome, America/Noronha, America/Panama, America/Pangnirtung, America/Paramaribo, America/Phoenix, America/Port-au-Prince, America/Port_of_Spain, America/Porto_Acre, America/Puerto_Rico, America/Rainy_River, America/Rankin_Inlet, America/Regina, America/Rosario, America/Santiago,

America/Santo_Domingo, America/Sao_Paulo, America/Scoresbysund, America/Shiprock, America/St_Johns, America/St_Kitts, America/St_Lucia, America/St_Thomas, America/St_Vincent, America/Swift_Current, America/Tegucigalpa, America/Thule, America/Thunder_Bay, America/Tijuana, America/Tortola, America/Vancouver, America/Whitehorse, America/Winnipeg, America/Yakutat, America/Yellowknife

Antarctica/Casey, Antarctica/DumontDUrville, Antarctica/Mawson, Antarctica/McMurdo, Antarctica/Palmer, Antarctica/South_Pole

Arctic/Longyearbyen

Asia/Aden, Asia/Alma-Ata, Asia/Amman, Asia/Anadyr, Asia/Aqtau, Asia/Aqtobe, Asia/Ashkhabad, Asia/Baghdad, Asia/Bahrain, Asia/Baku, Asia/Bangkok, Asia/Beirut, Asia/Bishkek, Asia/Brunei, Asia/Calcutta, Asia/Chungking, Asia/Colombo, Asia/Dacca, Asia/Damascus, Asia/Dubai, Asia/Dushanbe, Asia/Gaza, Asia/Harbin, Asia/Hong_Kong, Asia/Irkutsk, Asia/Ishigaki, Asia/Jakarta, Asia/Jayapura, Asia/Jerusalem, Asia/Kabul, Asia/Kamchatka, Asia/Karachi, Asia/Kashgar, Asia/Katmandu, Asia/Krasnoyarsk, Asia/Kuala_Lumpur, Asia/Kuching, Asia/Kuwait, Asia/Macao, Asia/Magadan, Asia/Manila, Asia/Muscat, Asia/Nicosia, Asia/Novosibirsk, Asia/Omsk, Asia/Phnom_Penh, Asia/Pyongyang, Asia/Qatar, Asia/Rangoon, Asia/Riyadh, Asia/Saigon, Asia/Seoul, Asia/Shanghai, Asia/Singapore, Asia/Taipei, Asia/Tashkent, Asia/Tbilisi, Asia/Tehran, Asia/Thimbu, Asia/Tokyo, Asia/Ujung_Pandang, Asia/Ulan_Bator, Asia/Urumqi, Asia/Vientiane, Asia/Vladivostok, Asia/Yakutsk, Asia/Yekaterinburg, Asia/Yerevan

Atlantic/Azores, Atlantic/Bermuda, Atlantic/Canary, Atlantic/Cape_Verde, Atlantic/Faeroe, Atlantic/Jan_Mayen, Atlantic/Madeira, Atlantic/Reykjavik, Atlantic/South_Georgia, Atlantic/St_Helena, Atlantic/Stanley

Australia/Adelaide, Australia/Brisbane, Australia/Broken_Hill, Australia/Darwin, Australia/Hobart, Australia/Lindeman, Australia/Lord_Howe, Australia/Melbourne, Australia/Perth, Australia/Sydney

Europe/Amsterdam, Europe/Andorra, Europe/Athens, Europe/Belfast, Europe/Belgrade, Europe/Berlin, Europe/Bratislava, Europe/Brussels, Europe/Bucharest, Europe/Budapest, Europe/Chisinau, Europe/Copenhagen, Europe/Dublin, Europe/Gibraltar, Europe/Helsinki, Europe/Istanbul, Europe/Kaliningrad, Europe/Kiev, Europe/Lisbon, Europe/Ljubljana, Europe/London, Europe/Luxembourg, Europe/Madrid, Europe/Malta, Europe/Minsk, Europe/Monaco, Europe/Moscow, Europe/Oslo, Europe/Paris, Europe/Prague, Europe/Riga, Europe/Rome, Europe/Samara, Europe/San_Marino, Europe/Sarajevo, Europe/Simferopol, Europe/Skopje, Europe/Sofia, Europe/Stockholm, Europe/Tallinn, Europe/Tirane, Europe/Vaduz, Europe/Vatican, Europe/Vienna, Europe/Vilnius, Europe/Warsaw, Europe/Zagreb, Europe/Zurich

Indian/Antananarivo, Indian/Chagos, Indian/Christmas, Indian/Cocos, Indian/Comoro, Indian/Kerguelen, Indian/Mahe, Indian/Maldives, Indian/Mauritius, Indian/Mayotte, Indian/Reunion

Pacific/Apia, Pacific/Auckland, Pacific/Chatham, Pacific/Easter, Pacific/Efate, Pacific/Enderbury, Pacific/Fakaofo, Pacific/Fiji, Pacific/Funafuti, Pacific/Galapagos, Pacific/Gambier, Pacific/Guadalcanal, Pacific/Guam, Pacific/Honolulu, Pacific/Johnston, Pacific/Kiritimati, Pacific/Kosrae, Pacific/Kwajalein, Pacific/Majuro, Pacific/Marquesas, Pacific/Midway, Pacific/Nauru, Pacific/Niue, Pacific/Norfolk, Pacific/Noumea, Pacific/Pago_Pago, Pacific/Palau, Pacific/Pitcairn, Pacific/Ponape, Pacific/Port_Moresby, Pacific/Rarotonga, Pacific/Saipan, Pacific/Tahiti, Pacific/Tarawa, Pacific/Tongatapu, Pacific/Truk, Pacific/Wake, Pacific/Wallis, Pacific/Yap

Required Privilege Level

system—To view this statement in the configuration.

system-control—To add this statement to the configuration.

Related Documentation

- Modifying the Default Time Zone for a Router or Switch Running Junos OS
- System Management Configuration Statements

traceoptions

Syntax

```
traceoptions {
  file <filename> <files number> <size size> <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
}
```

Hierarchy Level [edit **event-options event-script**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define tracing operations for event scripts.

Default If you do not include this statement, no event script–specific tracing operations are performed.

Options ***filename***—Name of the file to receive the output of the tracing operation. All files are placed in the directory **/var/log**. By default, event script process tracing output is placed in the file **escript.log**. If you include the **file** statement, you must specify a filename. To retain the default, you can specify **escript.log** as the filename.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed and compressed to ***trace-file*.0.gz**, then ***trace-file*.1.gz**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

Range: 2 through 1000

Default: 10 files

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **all**—Log all operations
- **events**—Log important events
- **input**—Log event script input data
- **offline**—Generate data for offline development
- **output**—Log event script output data
- **rpc**—Log event script RPCs
- **xslt**—Log the XSLT library

no-world-readable—Restrict file access to owner. This is the default.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed

and compressed to ***trace-file*.0.gz**. When ***trace-file*** again reaches its maximum size, ***trace-file*.0.gz** is renamed ***trace-file*.1.gz** and ***trace-file*** is renamed and compressed to ***trace-file*.0.gz**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and a filename.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10 KB through 1 GB

Default: 128 KB

world-readable—Enable unrestricted file access.

Required Privilege Level

maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation

- Tracing Event Script Processing

traceoptions

Syntax

```
traceoptions {
  file <filename> <files number> <match regular-expression> <size size> <world-readable |
    no-world-readable>;
  flag flag;
  no-remote-trace;
}
```

Hierarchy Level [edit **event-options**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define tracing operations for event policies.

Default If you do not include this statement, no event-policy-specific tracing operations are performed.

Options ***filename***—Name of the file to receive the output of the tracing operation. All files are placed in the directory **/var/log**. By default, commit script process tracing output is placed in the file **eventd**. If you include the **file** statement, you must specify a filename. To retain the default, you can specify **eventd** as the filename.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed and compressed to ***trace-file*.0.gz**, then ***trace-file*.1.gz**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

Range: 2 through 1000

Default: 3 files

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **all**—Log all operations
- **configuration**—Log reading of configuration at the **[edit event-options]** hierarchy level
- **events**—Log eventd processing
- **database**—Log events involving storage and retrieval in events database
- **server**—Log communication with processes that are generating events
- **timer-events**—Log internally generated events

match *regular-expression*—(Optional) Refine the output to include lines that contain the regular expression.

no-world-readable—Restrict file access to owner. This is the default.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed and compressed to ***trace-file*.0.gz**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0.gz** is renamed ***trace-file*.1.gz** and ***trace-file*** is renamed and compressed to ***trace-file*.0.gz**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and filename.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10 KB through 1 GB

Default: 128 KB

world-readable—(Optional) Enable unrestricted file access.

Required Privilege Level

maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation

- Tracing Event Policy Processing

traceoptions (Commit and Op Scripts)

Syntax

```
traceoptions {
  file <filename> <files number> <size size> <world-readable | no-world-readable>;
  flag flag;
  no-remote-trace;
}
```

Hierarchy Level [edit system scripts commit],
[edit system scripts op]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define tracing operations for commit or op scripts.

Default If you do not include this statement, no script-specific tracing operations are performed.

Options ***filename***—Name of the file to receive the output of the tracing operation. All files are placed in the directory **/var/log**. By default, commit script process tracing output is placed in the file **cscript.log** and op script process tracing is placed in the file **op-script.log**. If you include the **file** statement, you must specify a filename. To retain the default, you can specify **cscript.log** or **op-script.log** as the filename.

files ***number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed and compressed to ***trace-file*.0.gz**, then ***trace-file*.1.gz**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

Range: 2 through 1000

Default: 10 files

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **all**—Log all operations
- **events**—Log important events
- **input**—Log script input data
- **offline**—Generate data for offline development
- **output**—Log script output data
- **rpc**—Log script RPCs
- **xslt**—Log the XSLT library

no-world-readable—Restrict file access to owner. This is the default.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed and compressed to ***trace-file*.0.gz**. When ***trace-file*** again reaches its maximum size, ***trace-file*.0.gz** is renamed ***trace-file*.1.gz** and ***trace-file*** is renamed and compressed to ***trace-file*.0.gz**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and a filename.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10 KB through 1 GB

Default: 128 KB

world-readable—Enable unrestricted file access.

Required Privilege Level

maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation

- Tracing Commit Script Processing
- Tracing Op Script Processing

transfer-delay

Syntax transfer-delay *seconds*;

Hierarchy Level [edit **event-options destinations** *destination-name*],
[edit **event-options policy** *policy-name* **then event-script** *filename* **destination** *destination-name*],
[edit **event-options policy** *policy-name* **then execute-commands destination** *destination-name*],
[edit **event-options policy** *policy-name* **then upload** filename (*filename* | committed) destination *destination-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a delay before transferring files. This allows the files to be completely generated before the upload starts. If you configure a transfer delay at the **[edit event-options destination *destination-name*]** hierarchy level and at one of the **[edit event-options policy *policy-name* then ...]** hierarchy levels, the resulting delay is the sum of the two delays.

Default If you do not include this statement, there is no transfer delay.

Options ***seconds***—Duration of the delay before files are uploaded.

Required Privilege Level maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation

- Defining Destinations for File Archiving by Event Policies
- Configuring the Delay Before Files Are Uploaded by an Event Policy

trigger

Syntax trigger (on | after | until) *event-count*;

Hierarchy Level [edit **event-options policy** *policy-name* **within** *seconds*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure an event policy to be triggered if an event or set of events occurs ***event-count*** times within a specified time period.

Default If you do not include this statement, the policy is executed on receipt of the first configured event.

Options

after *event-count*—The policy is executed when the number of matching events received equals ***event-count*** + 1.

on *event-count*—The policy is executed when the number of matching events received equals ***event-count***.

until *event-count*—The policy is executed each time a matching event is received and stops being executed when the number of matching events received equals ***event-count***.

Required Privilege Level maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation

- Triggering an Event Policy Based on Event Count

upload

Syntax

```
upload filename  (filename | committed) destination destination-name {
    retry-count count retry-interval seconds;
    transfer-delay seconds;
    user-name username;
}
```

Hierarchy Level [edit **event-options policy** *policy-name* **then**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description On receipt of an event, upload the committed configuration file to a destination.

Options

destination ***destination-name***—Name of the destination for the uploaded file. It must be defined in the **destinations** statement at the **[edit event-options]** hierarchy level.

filename (*filename* | committed)—Name of the file to upload. Specify either the word **committed** to upload the most recently committed configuration file, or the filename of another file.

The remaining statements are explained separately.

Required Privilege Level maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation

- **destinations on page 566**
- Configuring an Event Policy to Upload Files

user (System Logging)

Syntax

```
user (username | *) {
    facility severity;
    match "regular-expression";
}
```

Hierarchy Level [edit system **syslog**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the logging of system messages to user terminals.

Options

***** (the asterisk)—Log messages to the terminal sessions of all users who are currently logged in.

facility—Class of messages to log. To specify multiple classes, include multiple ***facility severity*** statements. For a list of the facilities, see Junos OS System Logging Facilities and Message Severity Levels.

severity—Severity of the messages that belong to the facility specified by the paired ***facility*** name. Messages with severities the specified level and higher are logged. For a list of the severities, see Junos OS System Logging Facilities and Message Severity Levels.

username—Junos OS login name of the user whose terminal session is to receive system log messages. To log messages to more than one user's terminal session, include more than one **user** statement.

The remaining statement is explained separately.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Directing System Log Messages to a User Terminal.
- Junos OS System Logging Facilities and Message Severity Levels
- *Junos OS System Log Messages Reference*

user-name

Syntax user-name *username*;

Hierarchy Level [edit **event-options policy** *policy-name* **then event-script** *filename*],
[edit **event-options policy** *policy-name* **then execute-commands**],
[edit **event-options policy** *policy-name* **then upload** filename (*filename* | committed) destination *destination-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Associate a user with an action in an event policy. The event policy action is executed under the privileges of the associated user.

Default If you do not associate a user with an action, the action is executed as user **root**.

Options ***username***—A username that is configured at the **[edit system login]** hierarchy level.

Required Privilege Level maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation
- Changing the User Privilege Level for an Event Policy Action

within

Syntax
```
within seconds {
  events [ events ];
  not events [ events ];
  trigger (after | on | until) event-count;
}
```

Hierarchy Level [edit **event-options policy** *policy-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Create a list of events that must (or must not) occur within a specified time interval for the policy to be triggered.

The statements are explained separately.

Options ***seconds***—Interval between events.

Range: 60 through 604,800 seconds

Required Privilege Level maintenance—To view this statement in the configuration.
maintenance-control—To add this statement to the configuration.

Related Documentation
- Using Correlated Events to Trigger an Event Policy

world-readable

Syntax world-readable | no-world-readable;

Hierarchy Level [edit system syslog **archive**],
[edit system syslog file *filename* archive]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Grant all users permission to read log files, or restrict the permission only to the **root** user and users who have the Junos OS **maintenance** permission.

Default **no-world-readable**

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Specifying Log File Size, Number, and Archiving Properties
- *Junos OS System Log Messages Reference*

CHAPTER 42

Operational Mode Commands for System Monitoring

clear log

Syntax clear log *filename*
<all>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Remove contents of a log file.

Options *filename*—Name of the specific log file to truncate.

all—(Optional) Truncate the specified log file and delete all archived versions of it.

Required Privilege Level clear

Related Documentation
- **show log on page 657**

List of Sample Output **clear log on page 612**

Output Fields See **file list** for an explanation of output fields.

clear log The following sample commands list log file information, clear the contents of a log file, and then display the updated log file information:

```
user@host> file list lcc0-re0:/var/log/sampled detail
lcc0-re0:
--------------------------------------------------------------------------
-rw-r-----  1 root  wheel      26450 Jun 23 18:47 /var/log/sampled
total 1

user@host> clear log lcc0-re0:sampled
lcc0-re0:
--------------------------------------------------------------------------

user@host> file list lcc0-re0:/var/log/sampled detail
lcc0-re0:
--------------------------------------------------------------------------
-rw-r-----  1 root  wheel         57 Sep 15 03:44 /var/log/sampled
total 1
```

file archive

Syntax file archive destination *destination* source *source* <compress>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Archive, and optionally compress, one or multiple local system files as a single file, locally or at a remote location.

Options

destination *destination*—Destination of the archived file or files. Specify the destination as a URL or filename. The Junos OS adds one of the following suffixes if the destination filename does not already have it:

- For archived files—The suffix **.tar**
- For archived and compressed files—The suffix **.tgz**

source *source*—Source of the original file or files. Specify the source as a URL or filename.

compress—(Optional) Compress the archived file with the GNU zip (gzip) compression utility. The compressed files have the suffix **.tgz**.

Required Privilege Level maintenance

List of Sample Output

file archive (Multiple Files) on page 613
file archive (Single File) on page 613
file archive (with Compression) on page 614

Output Fields When you enter this command, you are provided feedback on the status of your request.

file archive (Multiple Files) The following sample command archives all message files in the local directory **/var/log/messages** as the single file **messages-archive.tar** in the same directory:

```
user@host> file archive source /var/log/messages* destination /var/log/messages-archive.tar
/usr/bin/tar: Removing leading / from absolute path names in the archive.

user@host>
```

file archive (Single File) The following sample command archives one message file in the local directory **/var/log/messages** as the single file **messages-archive.tar** in the same directory:

```
user@host> file archive source /var/log/messages destination /var/log/messages-archive.tar
/usr/bin/tar: Removing leading / from absolute path names in the archive.
user@host
```

file archive (with Compression)

The following sample command archives and compresses all message files in the local directory **/var/log/messages** as the single file **messages-archive.tgz** in the same directory:

```
user@host> file archive compress source /var/log/messages* destination
/var/log/messages-archive.tgz
/usr/bin/tar: Removing leading / from absolute path names in the archive.

user@host>
```

file checksum md5

Syntax file checksum md5 <*pathname*> *filename*

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Calculate the Message Digest 5 (MD5) checksum of a file.

Options *pathname*—(Optional) Path to a filename.

filename—Name of a local file for which to calculate the MD5 checksum.

Required Privilege Level maintenance

Related Documentation

- Configuring Checksum Hashes for a Commit Script in the *Junos OS Configuration and Diagnostic Automation Guide*
- Configuring Checksum Hashes for an Event Script in the *Junos OS Configuration and Diagnostic Automation Guide*
- Configuring Checksum Hashes for an Op Script in the *Junos OS Configuration and Diagnostic Automation Guide*
- Executing an Op Script from a Remote Site in the *JUNO Configuration and Diagnostic Automation Guide*
- **file checksum sha-256 on page 366**
- **file checksum sha1 on page 365**
- **op on page 207**

List of Sample Output **file checksum md5 on page 615**

Output Fields When you enter this command, you are provided feedback on the status of your request.

file checksum md5

```
user@host> file checksum md5 jbundle-5.3R2.4-export-signed.tgz
MD5 (jbundle-5.3R2.4-export-signed.tgz) = 2a3b69e43f9bd4893729cc16f505a0f5
```

file checksum sha1

Syntax file checksum sha1 <*pathname*> *filename*

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Calculate the Secure Hash Algorithm (SHA-1) checksum of a file.

Options *pathname*—(Optional) Path to a filename.

filename—Name of a local file for which to calculate the SHA-1 checksum.

Required Privilege Level maintenance

Related Documentation

- Configuring Checksum Hashes for a Commit Script in the *Junos OS Configuration and Diagnostic Automation Guide*
- Configuring Checksum Hashes for an Event Script in the *Junos OS Configuration and Diagnostic Automation Guide*
- Configuring Checksum Hashes for an Op Script in the *Junos OS Configuration and Diagnostic Automation Guide*
- Executing an Op Script from a Remote Site in the *Junos OS Configuration and Diagnostic Automation Guide*
- **file checksum md5 on page 364**
- **file checksum sha-256 on page 366**
- **op on page 207**

List of Sample Output **file checksum sha1 on page 616**

Output Fields When you enter this command, you are provided feedback on the status of your request.

file checksum sha1

```
user@host> file checksum sha1 /var/db/scripts/opscript.slax

SHA1 (/var/db/scripts/commitscript.slax) = ba9e47120c7ce55cff29afd73eacd370e162c676
```

file checksum sha-256

Syntax file checksum sha-256 <*pathname*> *filename*

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Calculate the Secure Hash Algorithm 2 family (SHA-256) checksum of a file.

Options *pathname*—(Optional) Path to a filename.

filename—Name of a local file for which to calculate the SHA-256 checksum.

Required Privilege Level maintenance

Related Documentation

- Configuring Checksum Hashes for a Commit Script in the *Junos OS Configuration and Diagnostic Automation Guide*
- Configuring Checksum Hashes for an Event Script in the *Junos OS Configuration and Diagnostic Automation Guide*
- Configuring Checksum Hashes for an Op Script in the *Junos OS Configuration and Diagnostic Automation Guide*
- Executing an Op Script from a Remote Site in the *Junos OS Configuration and Diagnostic Automation Guide*
- **file checksum md5 on page 364**
- **file checksum sha1 on page 365**
- **op on page 207**

List of Sample Output **file checksum sha-256 on page 617**

Output Fields When you enter this command, you are provided feedback on the status of your request.

file checksum sha-256

```
user@host> file checksum sha-256 /var/db/scripts/commitscript.slax

SHA256 (/var/db/scripts/commitscript.slax) =
94c2b061fb55399e15babd2529453815601a602b5c98e5c12ed929c9d343dd71
```

file compare

Syntax file compare (files *filename filename*)
< context | unified>
<ignore-white-space >

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Compare two local files and describe the differences between them in default, context, or unified output styles:

- **Default**—In the first line of output, **c** means lines were changed between the two files, **d** means lines were deleted between the two files, and **a** means lines were added between the two files. The numbers preceding this alphabetical marker represent the first file, and the lines after the alphabetical marker represent the second file. A left angle bracket (**<**) in front of output lines refers to the first file. A right angle bracket (**>**) in front of output lines refers to the second file.
- **Context**—Display is divided into two parts. The first part is the first file; the second part is the second file. Output lines preceded by an exclamation point (**!**) have changed. Additions are marked with a plus sign (**+**), and deletions are marked with a minus sign (**-**).
- **Unified**—Display is preceded by the line number from the first and the second file (*xx*,*xxx*,*x*). Before the line number, additions to the file are marked with a plus sign (**+**), and deletions to the file are marked with a minus sign (**-**). The body of the output contains the affected lines. Changes are viewed as additions plus deletions.

Options files *filename*—Names of two local files to compare.

context—(Optional) Display output in context format.

ignore-white-space—(Optional) Ignore changes in amount of white space.

unified—(Optional) Display output in unified format.

Required Privilege Level none

List of Sample Output
file compare files on page 619
file compare files context on page 619
file compare files unified on page 619
file compare files unified ignore-white-space on page 620

Output Fields When you enter this command, you are provided feedback on the status of your request.

file compare files

```
user@host> file compare files /tmp/one /tmp/two
100c100
<             full-name "File 1";
---
>             full-name "File 2";
102c102
<             class foo; # 'foo' is not defined
---
>             class super-user;
```

file compare files context

```
user@host> file compare files /tmp/one /tmp/two context
*** /tmp/one    Wed Dec  3 17:12:50 2003
--- /tmp/two    Wed Dec  3 09:13:14 2003
***************
*** 97,104 ****
                }
            }
            user bill {
!               full-name "Bill Smith";
!               class foo; # 'foo' is not defined
                authentication {
                    encrypted-password SECRET;
                }
--- 97,105 ----
                }
            }
            user bill {
!               full-name "Bill Smith";
!               uid 1089;
!               class super-user;
                authentication {
                    encrypted-password SECRET;
                }
```

file compare files unified

```
user@host> file compare files /tmp/one /tmp/two unified
--- /tmp/one    Wed Dec  3 17:12:50 2003
+++ /tmp/two    Wed Dec  3 09:13:14 2003
@@ -97,8 +97,9 @@
                }
            }
            user bill {
-               full-name "Bill Smith";
-               class foo; # 'foo' is not defined
+               full-name "Bill Smith";
+               uid 1089;
+               class super-user;
                authentication {
                    encrypted-passwordSECRET;
                }
```

file compare files unified ignore-white-space

```
user@host> file compare files /tmp/one /tmp/two unified ignore-white-space
--- /tmp/one    Wed Dec  3 09:13:10 2003
+++ /tmp/two    Wed Dec  3 09:13:14 2003
@@ -99,7 +99,7 @@
         user bill {
             full-name "Bill Smith";
             uid 1089;
-            class foo; # 'foo' is not defined
+            class super-user;
             authentication {
                  encrypted-password <SECRET>; # SECRET-DATA
             }
```

file copy

Syntax

file copy *source destination*
<source-address *address*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Copy files from one place to another on the local router or switch or between the local router or switch and a remote system.

Options

source—Source of the original file. Specify this as a URL or filename.

destination—Destination of the copied file. Specify this as a URL or filename. If you are copying a file to the current directory (your home directory on the local router or switch) and are not renaming the file, specify the destination with a period (.).

source-address *address*—(Optional) Source IP host address. This option is useful for specifying the source address of a secure copy (scp) file transfer.

Required Privilege Level

maintenance

List of Sample Output

file copy (A File from the Router to a PC) on page 621
file copy (A Configuration File Between Routing Engines) on page 621
file copy (A Log File Between Routing Engines) on page 621

Output Fields

When you enter this command, you are provided feedback on the status of your request.

file copy (A File from the Router to a PC)

```
user@host> file copy /var/tmp/rpd.core.4 berry:/c/junipero/tmp

...transferring.file..... |          0 KB |   0.3 kB/s | ETA: 00:00:00 | 100%
```

file copy (A Configuration File Between Routing Engines)

The following sample command copies a configuration file from Routing Engine 0 to Routing Engine 1:

```
user@host>  file copy /config/juniper.conf re1:/var/tmp/copied-juniper.conf
```

file copy (A Log File Between Routing Engines)

The following sample command copies a log file from Routing Engine 0 to Routing Engine 1:

```
user@host> file copy lcc0-re0:/var/log/chassisd lcc0-re1:/var/tmp
```

file delete

Syntax
file delete *filename*
<purge>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Delete a file on the local router or switch.

Options
filename—Name of the file to delete. For a routing matrix, include chassis information in the filename if the file to be deleted is not local to the Routing Engine from which the command is issued.

purge—(Optional) Overwrite regular files before deleting them.

Required Privilege Level
maintenance

List of Sample Output
file delete on page 622

Output Fields
When you enter this command, you are provided feedback on the status of your request.

file delete

```
user@host> file list /var/tmp
dcd.core
rpd.core
snmpd.core

user@host> file delete /var/tmp/snmpd.core
user@host> file list /var/tmp
dcd.core
rpd.core
```

file list

Syntax
file list
<detail | recursive>
<*filename*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display a list of files on the local router or switch.

Options
none—Display a list of all files for the current directory.

detail | recursive—(Optional) Display detailed output or descend recursively through the directory hierarchy, respectively.

filename—(Optional) Display a list of files. For a routing matrix, the filename must include the chassis information.

Additional Information
The default directory is the home directory of the user logged into the router or switch. To view available directories, enter a space and then a backslash (/) after the **file list** command. To view files within a specific directory, include a backslash followed by the directory and, optionally, subdirectory name after the **file list** command.

Required Privilege Level
maintenance

List of Sample Output
file list on page 623

Output Fields
When you enter this command, you are provided feedback on the status of your request.

file list

```
user@host> file list /var/tmp
dcd.core
rpd.core
snmpd.core
```

file rename

Syntax file rename *source destination*

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Rename a file on the local router or switch.

Options *destination*—New name for the file.

source—Original name of the file. For a routing matrix, the filename must include the chassis information.

Required Privilege Level maintenance

List of Sample Output **file rename on page 624**

Output Fields When you enter this command, you are provided feedback on the status of your request.

file rename The following example lists the files in **/var/tmp**, renames one of the files, and then displays the list of files again to reveal the newly named file.

```
user@host> file list /var/tmp
dcd.core
rpd.core
snmpd.core

user@host> file rename /var/tmp/dcd.core /var/tmp/dcd.core.990413
user@host> file list /var/tmp
dcd.core.990413
rpd.core
snmpd.core
```

file show

Syntax file show *filename*
<encoding base64>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display the contents of a file.

Options *filename*—Name of a file. For a routing matrix, the filename must include the chassis information.

encoding base64—(Optional) Encode file contents.

Required Privilege Level maintenance

List of Sample Output file show on page 625

Output Fields When you enter this command, you are provided feedback on the status of your request.

file show

```
user@host> file show /var/log/messages
Apr 13 21:00:08 romney /kernel: so-1/1/2: loopback suspected; going to standby.
Apr 13 21:00:40 romney /kernel: so-1/1/2: loopback suspected; going to standby.
Apr 13 21:02:48 romney last message repeated 4 times
Apr 13 21:07:04 romney last message repeated 8 times
Apr 13 21:07:13 romney /kernel: so-1/1/0: Clearing SONET alarm(s) RDI-P
Apr 13 21:07:29 romney /kernel: so-1/1/0: Asserting SONET alarm(s) RDI-P
...
```

monitor list

Syntax monitor list

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display the status of monitored log and trace files.

Options This command has no options.

Additional Information Log files are generated by the routing protocol process or by system logging. The log files generated by system logging are configured with the **syslog** statement at the **[edit system]** hierarchy level and the **options** statement at the **[edit routing-options]** hierarchy level. The trace files generated by the routing protocol process are those configured with **traceoptions** statements at the **[edit routing-options]**, **[edit interfaces]**, and **[edit protocols *protocol*]** hierarchy levels.

Required Privilege Level trace

Related Documentation
- **monitor start on page 627**
- **monitor stop on page 628**

List of Sample Output **monitor list on page 626**

Output Fields Table 82 on page 626 describes the output fields for the **monitor list** command. Output fields are listed in the approximate order in which they appear.

Table 82: monitor list Output Fields

Field Name	Field Description
monitor start	Indicates the file is being monitored.
"*filename*"	Name of the file that is being monitored.
Last changed	Date and time at which the file was last modified.

monitor list

```
user@host> monitor list
monitor start "vrrpd" (Last changed Dec 03:11:06 20)
monitor start "cli-commands" (Last changed Nov 07:3)
```

monitor start

Syntax monitor start *filename*

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Start displaying the system log or trace file and additional entries being added to those files.

Options *filename*—Specific log or trace file.

Additional Information Log files are generated by the routing protocol process or by system logging. The log files generated by system logging are configured with the **syslog** statement at the **[edit system]** hierarchy level and the **options** statement at the **[edit routing-options]** hierarchy level. The trace files generated by the routing protocol process are configured with **traceoptions** statements at the **[edit routing-options]**, **[edit interfaces]**, and **[edit protocols *protocol*]** hierarchy levels.

Required Privilege Level trace

Related Documentation

- **monitor list on page 626**
- **monitor stop on page 628**

List of Sample Output **monitor start on page 627**

Output Fields Table 83 on page 627 describes the output fields for the **monitor start** command. Output fields are listed in the approximate order in which they appear.

Table 83: monitor start Output Fields

Field Name	Field Description
******filename******	Name of the file from which entries are being displayed. This line is displayed initially and when the command switches between log files.
Date and time	Timestamp for the log entry.

monitor start

```
user@host> monitor start system-log
*** system-log***
Jul 20 15:07:34 hang sshd[5845]: log: Generating 768 bit RSA key.
Jul 20 15:07:35 hang sshd[5845]: log: RSA key generation complete.
Jul 20 15:07:35 hang sshd[5845]: log: Connection from 204.69.248.180 port 912
Jul 20 15:07:37 hang sshd[5845]: log: RSA authentication for root accepted.
Jul 20 15:07:37 hang sshd[5845]: log: ROOT LOGIN as 'root' from trip.jcmax.com
Jul 20 15:07:37 hang sshd[5845]: log: Closing connection to 204.69.248.180
```

monitor stop

Syntax monitor stop *filename*

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Stop displaying the system log or trace file.

Options *filename*—Specific log or trace file.

Additional Information Log files are generated by the routing protocol process or by system logging. The log files generated by system logging are those configured with the **syslog** statement at the **[edit system]** hierarchy level and the **options** statement at the **[edit routing-options]** hierarchy level. The trace files generated by the routing protocol process are those configured with **traceoptions** statements at the **[edit routing-options]**, **[edit interfaces]**, and **[edit protocols *protocol*]** hierarchy levels.

Required Privilege Level trace

Related Documentation

- **monitor list on page 626**
- **monitor start on page 627**

List of Sample Output **monitor stop on page 628**

Output Fields This command produces no output.

monitor stop

```
user@host> monitor stop
```

request system configuration rescue delete

Syntax request system configuration rescue delete

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Delete an existing rescue configuration.

Options This command has no options.

Required Privilege Level maintenance

Related Documentation

- request system configuration rescue save on page 377
- request system software rollback on page 113
- show system commit on page 381

List of Sample Output request system configuration rescue delete on page 629

Output Fields This command produces no output.

request system configuration rescue delete

```
user@host> request system configuration rescue delete
```

request system configuration rescue save

Syntax request system configuration rescue save

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Save the most recently committed configuration as the rescue configuration so that you can return to it at any time by using the **rollback** command.

Options This command has no options.

Required Privilege Level maintenance

Related Documentation

- **request system software delete on page 111**
- **request system software rollback on page 113**
- **show system commit on page 381**

List of Sample Output **request system configuration rescue save on page 630**

Output Fields This command produces no output.

request system configuration rescue save

```
user@host> request system configuration rescue save
```

request system scripts refresh-from commit

Syntax **request system scripts refresh-from commit file *file-name* url *url-path***

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Automatically download the initial Junos OS configuration and a set of standard commit scripts during a Junos XML management protocol/NETCONF session when a switch is brought up for the first time.

The Junos XML management protocol equivalent for this operational mode command is:

```
<request-script-refresh-from>
   <type>commit</type>
   <file>file-name></file>
   <URL>URL</URL>
</request-script-refresh-from>
```

Options file *file-name*—Name of the file to be downloaded.

url *url-path*—URL of the file to be downloaded.

Required Privilege Level maintenance

Related Documentation

- Understanding Automatic Refreshing of Scripts on J-EX Series Switches on page 323
- *Junos OS Junos XML Management Protocol Guide* at http://www.juniper.net/techpubs/software/junos/
- *Junos OS NETCONF XML Management Protocol Guide* at http://www.juniper.net/techpubs/software/junos/

List of Sample Output **request system scripts refresh-from commit file config.txt url http://host1.juniper.net on page 631**

request system scripts refresh-from commit file config.txt url http://host1.juniper.net

```
user@switch> request system scripts refresh-from commit file config.txt url
http://host1.juniper.net
user@switch>
```

request system scripts refresh-from event

Syntax **request system scripts refresh-from event file *file-name* url *url-path***

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Automatically download the initial Junos OS configuration and a set of standard event scripts during a Junos XML management protocol/NETCONF session when a switch is brought up for the first time.

The Junos XML management protocol equivalent for this operational mode command is:

```
<request-script-refresh-from>
   <type>event</type>
   <file>file-name></file>
   <URL>URL</URL>
</request-script-refresh-from>
```

Options file *file-name*—Name of the file to be downloaded.

url *url-path*—URL of the file to be downloaded.

Required Privilege Level maintenance

Related Documentation

- Understanding Automatic Refreshing of Scripts on J-EX Series Switches on page 323
- *Junos OS Junos XML Management Protocol Guide* at http://www.juniper.net/techpubs/software/junos/
- *Junos OS NETCONF XML Management Protocol Guide* at http://www.juniper.net/techpubs/software/junos/

List of Sample Output **request system scripts refresh-from event file config.txt url http://host1.juniper.net on page 632**

request system scripts refresh-from event file config.txt url http://host1.juniper.net

```
user@switch> request system scripts refresh-from event file config.txt url http://host1.juniper.net
user@switch>
```

request system scripts refresh-from op

Syntax **request system scripts refresh-from op file *file-name* url *url-path***

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Automatically download the initial Junos OS configuration and a set of standard op scripts during a Junos XML management protocol/NETCONF session when a switch is brought up for the first time.

The Junos XML management protocol equivalent for this operational mode command is:

```
<request-script-refresh-from>
   <type>op</type>
   <file>file-name></file>
   <URL>URL</URL>
</request-script-refresh-from>
```

Options file *file-name*—Name of the file to be downloaded.

url *url-path*—URL of the file to be downloaded.

Required Privilege Level maintenance

Related Documentation

- Understanding Automatic Refreshing of Scripts on J-EX Series Switches on page 323
- *Junos OS Junos XML Management Protocol Guide* at http://www.juniper.net/techpubs/software/junos/
- *Junos OS NETCONF XML Management Protocol Guide* at http://www.juniper.net/techpubs/software/junos/

List of Sample Output **request system scripts refresh-from op file config.txt url http://host1.juniper.net on page 633**

request system scripts refresh-from op file config.txt url http://host1.juniper.net

```
user@switch> request system scripts refresh-from op file config.txt url http://host1.juniper.net
user@switch>
```

show chassis alarms

Syntax show chassis alarms

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display information about the conditions that have been configured to trigger alarms.

Options none—Display information about the conditions that have been configured to trigger alarms.

Additional Information You cannot clear the alarms for chassis components. Instead, you must remedy the cause of the alarm. When a chassis alarm is lit, it indicates that you are running the router or switch in a manner that we do not recommend.

Required Privilege Level view

List of Sample Output
show chassis alarms (Alarms Active) on page 634
show chassis alarms (No Alarms Active) on page 634
show chassis alarms (Backup Routing Engine) on page 634

Output Fields Table 84 on page 634 lists the output fields for the **show chassis alarms** command. Output fields are listed in the approximate order in which they appear.

Table 84: show chassis alarms Output Fields

Field Name	Field Description
Alarm time	Date and time the alarm was first recorded.
Class	Severity class for this alarm: **Minor** or **Major**.
Description	Information about the alarm.

show chassis alarms (Alarms Active)

```
user@host> show chassis alarms
3 alarms are currently active
Alarm time               Class  Description
2000-02-07 10:12:22 UTC Major fxp0: ethernet link down
2000-02-07 10:11:54 UTC Minor YELLOW ALARM - PEM 1 Removed
2000-02-07 10:11:03 UTC Minor YELLOW ALARM - Lower Fan Tray Removed
```

show chassis alarms (No Alarms Active)

```
user@host> show chassis alarms
No alarms are currently active
```

show chassis alarms (Backup Routing Engine)

```
user@host> show chassis alarms
2 alarms are currently active
Alarm time               Class  Description
2005-04-07 10:12:22 PDT  Minor  Host 1 Boot from alternate media
2005-04-07 10:11:54 PDT  Major  Host 1 compact-flash missing in Boot List
```

show chassis environment

Syntax show chassis environment

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display environmental information about the router or switch chassis, including the temperature and information about the fans, power supplies, and Routing Engine.

Options none—Display environmental information about the router or switch chassis.

For information about the remaining options, see the Related Topics.

Required Privilege Level view

Related Documentation

- show chassis environment cb
- show chassis environment cip
- **show chassis environment fpc on page 637**
- show chassis environment fpm
- show chassis environment mcs
- show chassis environment pcg
- show chassis environment pem
- **show chassis environment routing-engine on page 638**

List of Sample Output **show chassis environment (J-EX8208 Switch) on page 636**

Output Fields Table 85 on page 635 lists the output fields for the **show chassis environment** command. Output fields are listed in the approximate order in which they appear.

Table 85: show chassis environment Output Fields

Field Name	Field Description
Class	Item, Status, Measurement
Power	Information about each power supply. **Status** can be **OK**, **Testing** (during initial power-on), **Check**, **Failed**, or **Absent**.
Temp	Temperature of air flowing through the chassis. **Measurement** indicates degrees in Celsius (C) and Fahrenheit (F).
Fan	Information about the fans. **Status** can be **OK**, **Testing** (during initial power-on), **Failed**, or **Absent**. **Measurement** indicates if fans are spinning at normal or high speed.
Misc	Information about other components of the chassis.

show chassis environment (J-EX8208 Switch)

```
user@switch> show chassis environment
Class Item                           Status     Measurement
Power PSU 0                          OK
      PSU 1                          OK
      PSU 2                          OK
      PSU 3                          Check
      PSU 4                          Check
      PSU 5                          Check
Temp  CB 0 Intake                    OK         20 degrees C / 68 degrees
      CB 0 Exhaust                   OK         24 degrees C / 75 degrees
      CB 1 Intake                    OK         19 degrees C / 66 degrees
      CB 1 Exhaust                   OK         23 degrees C / 73 degrees
      CB 2 Intake                    OK         19 degrees C / 66 degrees
      CB 2 Exhaust                   OK         23 degrees C / 73 degrees
Fans  Fan 1                          OK         Spinning at normal speed
      Fan 2                          OK         Spinning at normal speed
      Fan 3                          OK         Spinning at normal speed
      Fan 4                          OK         Spinning at normal speed
      Fan 5                          OK         Spinning at normal speed
      Fan 6                          OK         Spinning at normal speed
      Fan 7                          OK         Spinning at normal speed
      Fan 8                          OK         Spinning at normal speed
      Fan 9                          OK         Spinning at normal speed
      Fan 10                         OK         Spinning at normal speed
      Fan 11                         OK         Spinning at normal speed
      Fan 12                         OK         Spinning at normal speed
```

show chassis environment fpc

Syntax show chassis environment fpc
<*slot*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description (J-EX Series switches) Display environmental information about Flexible PIC Concentrators (FPCs).

Options none—Display environmental information about all FPCs.

slot—(Optional) Display environmental information about an individual FPC:

- user@host> **show chassis environment fpc 1 lcc 1**
  user@host> **show chassis environment fpc 9**
- J-EX Series switches:
  - J-EX4200 standalone switches—Replace ***slot*** with 0.
  - J-EX4200 switches in a Virtual Chassis configuration—Replace ***slot*** with a value from 0 through 9 (switch's member ID).
  - J-EX8208 switches—Replace ***slot*** with a value from 0 through 7 (line card).
  - J-EX8216 switches—Replace ***slot*** with a value from 0 through 15 (line card).

Required Privilege Level view

Output Fields Table 86 on page 637 lists the output fields for the **show chassis environment fpc** command. Output fields are listed in the approximate order in which they appear.

Table 86: show chassis environment fpc Output Fields

Field Name	Field Description
State	Status of the FPC: • **Unknown**—FPC is not detected by the router. • **Empty**—No FPC is present. • **Present**—FPC is detected by the chassis daemon but is either not supported by the current version of the Junos OS, or the FPC is coming up but not yet online. • **Ready**—FPC is in intermediate or transition state. • **Announce online**—Intermediate state during which the FPC is coming up but not yet online, and the chassis manager acknowledges the chassisd FPC online initiative. • **Online**—FPC is online and running. • **Offline**—FPC is powered down. • **Diagnostics**—FPC is set to operate in diagnostics mode.
Power	Information about the voltage supplied to the FPC. The left column displays the required power, in volts. The right column displays the measured power, in millivolts.

show chassis environment routing-engine

Syntax show chassis environment routing-engine
<*slot*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display Routing Engine environmental status information.

Options none—Display environmental information about all Routing Engines.

slot—(Optional) Display environmental information about an individual Routing Engine. . On J-EX4200 standalone switches, replace ***slot*** with 0. On J-EX4200 switches in a Virtual Chassis configuration and on J-EX8208 and J-EX8216 switches, replace ***slot*** with 0 or 1.

Required Privilege Level view

List of Sample Output **show chassis environment routing-engine (Nonredundant) on page 638**
show chassis environment routing-engine (Redundant) on page 638

Output Fields Table 87 on page 638 lists the output fields for the **show chassis environment routing-engine** command. Output fields are listed in the approximate order in which they appear.

Table 87: show chassis environment routing-engine Output Fields

Field Name	Field Description
Routing engine *slot* status	Number of the Routing Engine slot: **0** or **1**.
State	Status of the Routing Engine: • **Online Master**—MCS is online, operating as Master. • **Online Standby**—MCS is online, operating as Standby.
Temperature	Temperature of the air flowing past the Routing Engine.

show chassis environment routing-engine (Nonredundant)

```
user@host> show chassis environment routing-engine
Routing Engine 0 status:
  State                          Online Master
  Temperature                 27 degrees C / 80 degrees
```

show chassis environment routing-engine (Redundant)

```
user@host> show chassis environment routing-engine
Route Engine 0 status:
  State:                       Online Master
  Temperature:               26 degrees C / 78 degrees F
Route Engine 1 status:
  State:                       Online Standby
  Temperature:               26 degrees C / 78 degrees F
```

show chassis fpc

Syntax
show chassis fpc
<detail <*fpc-slot*>> | <pic-status <*fpc-slot*>>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display status information about the installed Flexible PIC Concentrators (FPCs) and PICs.

Options
none—Display status information for all FPCs.

detail—(Optional) Display detailed status information for all FPCs or for the FPC in the specified slot (see ***fpc-slot***).

fpc-slot—(Optional) FPC slot number:

- J-EX Series switches:
  - J-EX4200 standalone switches—Replace ***fpc-slot*** with 0.
  - J-EX4200 switches in a Virtual Chassis configuration—Replace ***fpc-slot*** with a value from 0 through 9 (switch's member ID).
  - J-EX8208 switches—Replace ***fpc-slot*** with a value from 0 through 7 (line card).
  - J-EX8216 switches—Replace ***fpc-slot*** with a value from 0 through 15 (line card).

pic-status—(Optional) Display status information for all PICs or for the PIC in the specified slot (see ***fpc-slot***).

Required Privilege Level
view

Related Documentation
- request chassis fpc

List of Sample Output
show chassis fpc (Hardware Not Supported) on page 641
show chassis fpc detail (Hardware Not Supported) on page 641
show chassis fpc pic-status on page 641

Output Fields
Table 88 on page 640 lists the output fields for the**show chassis fpc** command. Output fields are listed in the approximate order in which they appear.

Table 88: show chassis fpc Output Fields

Field Name	Field Description	Level of Output
Slot or **Slot State**	Slot number and state. The state can be one of the following conditions: • **Dead**—Held in reset because of errors. • **Diag**—Slot is being ignored while the FPC is running diagnostics. • **Dormant**—Held in reset. • **Empty**—No FPC is present. • **Online**—FPC is online and running. • **Present**—FPC is detected by the chassis daemon but is either not supported by the current version of Junos OS or inserted in the wrong slot. The output also states either **Hardware Not Supported** or **Hardware Not In Right Slot**. FPC is coming up but not yet online. • **Probed**—Probe is complete; awaiting restart of the Packet Forwarding Engine (PFE). • **Probe-wait**—Waiting to be probed.	all levels
Logical slot	Slot number.	all levels
Temp (C) or Temperature	Temperature of the air passing by the FPC, in degrees Celsius or in both Celsius and Fahrenheit.	all levels
Total CPU Utilization (%)	Total percentage of CPU being used by the FPC's processor.	all levels
Interrupt CPU Utilization (%)	Of the total CPU being used by the FPC's processor, the percentage being used for interrupts.	none specified
Memory DRAM (MB)	Total DRAM, in megabytes, available to the FPC's processor.	none specified
Heap Utilization (%)	Percentage of heap space (dynamic memory) being used by the FPC's processor. If this number exceeds 80 percent, there may be a software problem (memory leak).	none specified
Buffer Utilization (%)	Percentage of buffer space being used by the FPC's processor for buffering internal messages.	none specified
Total CPU DRAM	Amount of DRAM available to the FPC's CPU.	**detail**
Total RLDRAM	Amount of reduced latency dynamic random access memory (RLDRAM) available to the FPC CPU.	**detail**
Total DDR DRAM	Amount of double data rate dynamic random access memory (DDR DRAM) available to the FPC CPU.	**detail**
Total SRAM	Amount of static RAM (SRAM) used by the FPC's CPU.	**detail**
Total SDRAM	Total amount of memory used for storing packets and notifications.	**detail**

Table 88: show chassis fpc Output Fields *(continued)*

Field Name	Field Description	Level of Output
I/O Manager ASICs information	I/O Manager version number, manufacturer, and part number.	**detail**
Start time	Time when the Routing Engine detected that the FPC was running.	**detail**
Uptime	How long the Routing Engine has been connected to the FPC and, therefore, how long the FPC has been up and running.	**detail**
PIC type	(**pic-status** output only) Type of PIC.	none specified

show chassis fpc (Hardware Not Supported)

```
user@host> show chassis fpc
show chassis fpc
                     Temp  CPU Utilization (%)   Memory    Utilization (%)
 Slot State            (C)  Total  Interrupt      DRAM (MB) Heap     Buffer
  0  Online          -------------------- CPU less FPC ---------------------
  1  Present         ------ Hardware Not In Right Slot ------
  2  Online                    0          0         0         0          0
  3  Present         ------ Hardware Not Supported ------
  4  Empty
  5  Empty
  6  Online                    0          0         0         0          0
```

show chassis fpc detail (Hardware Not Supported)

```
user@host> show chassis fpc detail
Slot 0 information:
  State                               Online
  Total CPU DRAM                      ---- CPU less FPC ----
  Start time                          2006-07-07 03:21:00 UTC
  Uptime                              27 minutes, 51 seconds
 Slot 1 information:
  State                               Present
  Reason                              --- Hardware Not In Right Slot ---
 Slot 2 information:
  State                               Online
  Total CPU DRAM                    32 MB
  Start time                          2006-07-07 03:20:59 UTC
  Uptime                              27 minutes, 52 seconds
 Slot 3 information:
  State                               Present
  Reason                              --- Hardware Not Supported ---
  Total CPU DRAM                     0 MB
 Slot 6 information:
  State                               Online
  Total CPU DRAM                    32 MB
  Start time                          2006-07-07 03:21:01 UTC
  Uptime                              27 minutes, 50 seconds
```

show chassis fpc pic-status

```
user@host> show chassis fpc pic-status
Slot 0 Online
  PIC 1    1x OC-12 ATM, MM
  PIC 2    1x OC-12 ATM, MM
  PIC 3    1x OC-12 ATM, MM
Slot 1 Online
```

```
  PIC 0    1x OC-48 SONET, SMIR
Slot 2 Online
  PIC 0    1x OC-192 SONET, SMSR
```

show chassis hardware

Syntax

```
show chassis hardware
<clei-models>
<detail | extensive>
<models>
```

Syntax (J-EX4200 Switch)

```
show chassis hardware
<detail | extensive>
```

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display a list of all Flexible PIC Concentrators (FPCs) and PICs installed in the router or switch chassis, including the hardware version level and serial number.

In J-EX Series switch command output, FPC refers to the following:

- On J-EX4200 standalone switches—Refers to the switch; ***fpc-number*** is always 0.
- On J-EX4200 switches in a Virtual Chassis configuration—Refers to the member of a Virtual Chassis; FPC ***number*** equals the member ID, from 0 through 9.
- On J-EX8208 and J-EX8216 switches—Refers to a line card; FPC ***number*** equals the slot number for the line card.

Options none—Display information about hardware.

clei-models—(Optional) Display Common Language Equipment Identifier (CLEI) bar code and model number for orderable field-replaceable units (FRUs).

detail—(Optional) Include RAM and disk information in output.

extensive—(Optional) Display ID EEPROM information.

models—(Optional) Display model numbers and part numbers for orderable FRUs and, for components that use ID EEPROM format v2, the CLEI code.

Required Privilege Level view

List of Sample Output
show chassis hardware (J-EX8216 Switch) on page 644
show chassis hardware clei-models (J-EX8216 Switch) on page 645

Output Fields Table 89 on page 644 lists the output fields for the **show chassis hardware** command. Output fields are listed in the approximate order in which they appear.

Table 89: show chassis hardware Output Fields

Field Name	Field Description	Level of Output
Item	Chassis—Information about the chassis, Routing Engine (SRE and RE modules in J-EX8200 switches), power supplies, fan trays, and LCD panel. Also displays information about Flexible PIC Concentrators (FPCs) and associated Physical Interface Cards (PICs). Information about the backplane, midplane, and SIBs (SF modules) is displayed for J-EX8200 switches. See J-EX Series Switches Hardware and CLI Terminology Mapping.	All levels
Version	Revision level of the chassis component.	All levels
Part number	Part number of the chassis component.	All levels
Serial number	Serial number of the chassis component. The serial number of the backplane is also the serial number of the chassis. Use this serial number when you need to contact Dell Support (see "Requesting Technical Support" on page lxxi) about the chassis.	All levels
Assb ID or **Assembly ID**	(**extensive** output only) Identification number that describes the FRU hardware.	All levels
FRU model number	(**clei-models**, **extensive**, and **models** keyword only) Model number of FRU hardware component.	none specified
CLEI code	(**clei-models** and **extensive**keyword only) Common Language Equipment Identifier code. This value is displayed only for hardware components that use ID EEPROM format v2. This value is not displayed for components that use ID EEPROM format v1.	none specified
EEPROM Version	ID EEPROM version used by hardware component: **0x01** (version 1) or **0x02** (version 2).	**extensive**
Description	Brief description of the hardware item: • Type of power supply. • Type of PIC. If the PIC type is not supported on the current software release, the output states **Hardware Not Supported** • Type of FPC: **FPC Type 1**, **FPC Type 2**, **FPC Type 3**, **FPC Type 4** , or **FPC Type OC192**, . A brief description of the FPC. • **MPC M 16x 10GE**—16-port 10-Gigabit Module Port Concentrator that supports SFP+ optical transceivers. (Not on J-EX Series switches.) • For hosts, the Routing Engine type. • For small form-factor pluggable transceiver (SFP) modules, the type of fiber: **LX**, **SX**, **LH**, or **T**. • LCD description for J-EX Series switches.	All levels

show chassis hardware (J-EX8216 Switch)

```
user@host> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis          REV 06                CY0109260072      DELL J-EX8216
Midplane         REV 06   710-016845   BA0909160167      EX8216-MP
```

```
CB 0             REV 22   710-020771   AX0109197708      EX8216-RE320
  Routing Engine 0        BUILTIN      BUILTIN           RE-EX8216
CB 1             REV 22   710-020771   AX0109197755      EX8216-RE320
  Routing Engine 1        BUILTIN      BUILTIN           RE-EX8216
FPC 5            REV 20   710-020683   BC0109228159      EX8200-48F
  CPU            REV 13   710-020598   BF0109197545      EX8200-CPU
  PIC 0                   BUILTIN      BUILTIN           48x 100 Base-FX/1000
Base-X
SIB 0            REV 10   710-021613   AY0109207864      EX8216-SF320
SIB 1            REV 10   710-021613   AY0109207808      EX8216-SF320
SIB 2            REV 10   710-021613   AY0109207917      EX8216-SF320
SIB 3            REV 10   710-021613   AY0109207831      EX8216-SF320
SIB 4            REV 10   710-021613   AY0109207811      EX8216-SF320
SIB 5            REV 10   710-021613   AY0109207881      EX8216-SF320
SIB 6            REV 10   710-021613   AY0109207837      EX8216-SF320
SIB 7            REV 10   710-021613   AY0109207819      EX8216-SF320
PSU 0            REV 01   740-030762   BG0709251730      EX8200-AC2K
PSU 1            REV 01   740-030762   BG0709251728      EX8200-AC2K
PSU 2            REV 01   740-030762   BG0709251743      EX8200-AC2K
PSU 3            REV 01   740-030762   BG0709251741      EX8200-AC2K
PSU 4            REV 01   740-030762   BG0709251729      EX8200-AC2K
PSU 5            REV 01   740-030762   BG0709251737      EX8200-AC2K
Top Fan Tray
  FTC 0          REV 1    760-030533   CX1209110149      EX8216-FT
  FTC 1          REV 1    760-030533   CX1209110149      EX8216-FT
Bottom Fan Tray
  FTC 0          REV 1    760-030533   CX1209110121      EX8216-FT
  FTC 1          REV 1    760-030533   CX1209110121      EX8216-FT
LCD 0            REV 04   710-025742   CE0109020194      EX8200 LCD
```

show chassis hardware clei-models (J-EX8216 Switch)

```
user@host> show chassis hardware clei-models
Hardware inventory:
Item             Version  Part number  CLEI code         FRU model number
Midplane         REV 08   710-016845
PSU 0            REV 05   740-023002   COUPAEAEAA        EX8200-PWR-AC3KR
PSU 1            REV 05   740-023002   COUPAEAEAA        EX8200-PWR-AC3KR
PSU 2            REV 05   740-023002   COUPAEAEAA        EX8200-PWR-AC3KR
PSU 3            REV 05   740-023002   COUPAEAEAA        EX8200-PWR-AC3KR
PSU 4            REV 05   740-023002   COUPAEAEAA        EX8200-PWR-AC3KR
PSU 5            REV 05   740-023002   COUPAEAEAA        EX8200-PWR-AC3KR
Top Fan Tray
Bottom Fan Tray
```

show chassis led

Syntax **show chassis led**
<fpc-slot <*fpc-slot-number*>>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display the status and colors of the chassis LEDs on the front of the switch. A major alarm (red) indicates a critical error condition that requires immediate action. A minor alarm (yellow) indicates a noncritical condition that requires monitoring or maintenance. A minor alarm that is left unchecked might cause interruption in service or performance degradation.

Options none—Display the status of the chassis status LEDs (for J-EX4200 switches configured as a Virtual Chassis, display the information for all Virtual Chassis members).

fpc-slot <*fpc-slot-number*>—(Optional) Display the information as follows:

- For the standalone J-EX4200 switch (***fpc-slot-number*** equals 0)
- For all J-EX4200 switches in a Virtual Chassis (**fpc-slot** with no ***fpc-slot-number*** value specified)
- For a specific Virtual Chassis member (***fpc-slot-number*** equals member ID value)
- For the line card in the specified slot on a J-EX8200 switch (***fpc-slot-number*** equals slot number)

Required Privilege Level view

Related Documentation
- Chassis Status LEDs in J-EX4200 Switches
- Chassis Status LEDs in a J-EX8200 Switch

Output Fields Table 90 on page 646 lists the output fields for the **show chassis led** command. Output fields are listed in the approximate order in which they appear.

Table 90: show chassis led Output Fields

Field Name	Field Description
Front panel contents for slot	FPC slot number of the switch whose content is being displayed. The number is always **0**, except for J-EX4200 switches in a Virtual Chassis, where it is the member ID value.
Front panel contents (EX8200 switches)	On J-EX8200 switches, no slot number is displayed.

Table 90: show chassis led Output Fields *(continued)*

Field Name	Field Description
Alarms LED	Displays status of the ALM LED: • Off—No alarm has been configured. • Green—No alarm has been triggered. • Red—Major alarm. • Yellow—Minor alarm
System LED	Displays status of the SYS LED: • Off—Switch is powered off. • Green—Switch is operating normally. • Yellow—Switch is booting.

Master LED	Displays status of the MST LED (on J-EX4200, and J-EX8200 switches):
	• Green—On a J-EX4200 Virtual Chassis switch, indicates the switch is the master in the Virtual Chassis configuration. On other switches, indicates that the Routing Engine is operational.
	• Off
	  ◦ On a J-EX4200 Virtual Chassis switch, indicates that this switch is not the master in the Virtual Chassis configuration.
	  ◦ On standalone J-EX4200, and J-EX8200 switches, indicates that the Routing Engine is not operational.

Interface	Names of the interfaces on the switch.
LED (ADM/SPD/DPX/POE)	State of the currently selected port parameter of the Status LED for the interface. The Status LED port parameters are: NOTE: J-EX8200 switches do not have the POE port parameter. • **ADM**—Administrative • **SPD**—Speed • **DPX**—Duplex • **POE**—Power over Ethernet

show chassis led

```
user@switch> show chassis led

Front panel contents for slot: 0
---------------------------------
LEDs status:
    Alarms LED: Off
    System LED: Green
    Master LED: Green
Interface       LED(ADM/SPD/DPX/POE)
-------------------------------------
ge-0/0/0        Off
ge-0/0/1        Full Duplex
ge-0/0/2        Full Duplex
ge-0/0/3        Off
ge-0/0/4        Off
```

```
ge-0/0/5        Full Duplex
ge-0/0/6        Full Duplex
ge-0/0/7        Full Duplex
ge-0/0/8        Full Duplex
ge-0/0/9        Full Duplex
ge-0/0/10       Full Duplex
ge-0/0/11       Full Duplex
ge-0/0/12       Full Duplex
ge-0/0/13       Full Duplex
ge-0/0/14       Full Duplex
ge-0/0/15       Full Duplex
ge-0/0/16       Full Duplex
ge-0/0/17       Full Duplex
ge-0/0/18       Full Duplex
ge-0/0/19       Full Duplex
ge-0/0/20       Full Duplex
ge-0/0/21       Full Duplex
ge-0/0/22       Off
ge-0/0/23       Off
ge-0/0/24       Full Duplex
ge-0/0/25       Full Duplex
ge-0/0/26       Off
ge-0/0/27       Off
ge-0/0/28       Full Duplex
ge-0/0/29       Full Duplex
```

show chassis led fpc-slot 0

```
user@switch> show chassis led fpc-slot 0
Front panel contents for slot: 0
---------------------------------
LEDs status:
    Alarms LED: Red
    System LED: Green
    Master LED: Green
Interface       LED(ADM/SPD/DPX/POE)
-------------------------------------
ge-0/0/0        Off
ge-0/0/1        Off
ge-0/0/2        Off
ge-0/0/3        Off
ge-0/0/4        Off
ge-0/0/5        Off
ge-0/0/6        Off
ge-0/0/7        Off
ge-0/0/8        Off
ge-0/0/9        Off
ge-0/0/10       Off
ge-0/0/11       Off
ge-0/0/12       Off
ge-0/0/13       Off
ge-0/0/14       Off
ge-0/0/15       Off
ge-0/0/16       Off
ge-0/0/17       Off
ge-0/0/18       Off
ge-0/0/19       Off
ge-0/0/20       Off
ge-0/0/21       Off
ge-0/0/22       Off
ge-0/0/23       Off
```

show chassis location

Syntax show chassis location

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display the physical location of the chassis. This command can only be used on the master Routing Engine.

Options none—Display all information about the physical location of the chassis.

Required Privilege Level view

List of Sample Output **show chassis location on page 649**
show chassis location on page 649

Output Fields Table 91 on page 649 lists the output fields for the **show chassis location** command. Output fields are listed in the approximate order in which they appear.

Table 91: show chassis location Output Fields

Field Name	Field Description
country-code	Country code information.
postal-code	Postal code information.
Building	Building information.
Floor	Floor information.
Global FPC	Global FPC number. The FPC slot number, when all FPC slots in the Routing Matrix are considered. The range of values is **0** through **31**.
LCC	Line-card chassis number.
Local FPC	Local FPC number.

show chassis location

```
user@host> show chassis location
country-code: US
postal-code: 94404
Building: Building 2, Floor: 2
```

show chassis location

```
user@host> show chassis location
country-code: US
postal-code: 94404
Building: Building 2, Floor: 2
```

show chassis pic

Syntax show chassis pic fpc-slot *slot-number* pic-slot *slot-number*

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display status information about the PIC installed in the specified Flexible PIC Concentrator (FPC) and PIC slot.

Options fpc-slot *slot-number*—Display information about the PIC in this particular FPC slot:

- J-EX Series switches:
  - J-EX4200 standalone switches—Replace ***slot-number*** with 0.
  - J-EX4200 switches in a Virtual Chassis configuration—Replace ***slot-number*** with a value from 0 through 9 (switch's member ID).
  - J-EX8208 switches—Replace ***slot-number*** with a value from 0 through 7 (line card).
  - J-EX8216 switches—Replace ***slot-number*** with a value from 0 through 15 (line card).

pic-slot *slot-number*—Display information about the PIC in this particular PIC slot. For routers, replace ***slot-number*** with a value from **0** through **3**. For J-EX4200 switches, replace ***slot-number*** with 0 for built-in network interfaces and 1 for interfaces on uplink modules. For J-EX8208 and J-EX8216 switches, replace ***slot-number*** with 0.

Required Privilege Level view

Related Documentation

- **request chassis pic on page 209**

List of Sample Output

show chassis pic fpc-slot pic-slot on page 651
show chassis pic fpc-slot pic-slot (PIC Offline) on page 651
show chassis pic fpc-slot pic-slot (FPC Offline) on page 652
show chassis pic fpc-slot pic-slot (FPC Not Present) on page 652
show chassis pic fpc-slot pic-slot (PIC Not Present) on page 652

Output Fields Table 92 on page 650 lists the output fields for the **show chassis pic** command. Output fields are listed in the approximate order in which they appear.

Table 92: show chassis pic Output Fields

Field Name	Field Description
Type	PIC type.
ASIC type	Type of ASIC on the PIC.

Table 92: show chassis pic Output Fields *(continued)*

Field Name	Field Description
State	Status of the PIC. State is displayed only when a PIC is in the slot. • **Online**— PIC is online and running. • **Offline**—PIC is powered down.
PIC version	PIC hardware version.
Uptime	How long the PIC has been online.
Package	(MultiServices PICs only) Services package supported: **Layer-2** or **Layer-3**.
PIC Port Information	Port-level information for the PIC.
Port Number	Port number for the PIC.
Cable Type	Type of cable connected to the port: **LH**, **LX**, or **SX**

PIC Port Information (MX960 Router Bidirectional Optics)	Port-level information for the PIC.
	• Port—Port number
	• Cable type—Type of small form-factor pluggable (SFP) optical transceiver installed. Uplink interfaces display -U. Down link interfaces display -D.
	• Fiber type—Type of fiber. SM is single-mode.
	• Xcvr vendor—Transceiver vendor name.
	• Xcvr vendor part number—Transceiver vendor part number.
	  ◦ BX10-10-km bidirectional optics.
	  ◦ BX40-40-km bidirectional optics.
	  ◦ SFP-LX-40-km SFP optics.
	• Wavelength—Wavelength of the transmitted signal. Uplinks are always 1310 nm. Downlinks are either 1490 nm or 1550 nm.

show chassis pic fpc-slot pic-slot

```
user@host> show chassis pic fpc-slot 2 pic-slot 0
PIC fpc slot 2 pic slot 0 information:
  Type                            10x 1GE(LAN), 1000 BASE
  ASIC type                       H chip
  State                           Online
  PIC version                     1.1
  Uptime                          1 day, 50 minutes, 58 seconds
PIC Port Information:
  Port        Cable
  Number      Type
  0           GIGE 1000LX
  6           GIGE 1000LX
```

show chassis pic fpc-slot pic-slot (PIC Offline)

```
user@host> show chassis pic fpc-slot 1 pic-slot 0
```

```
PIC fpc slot 1 pic slot 0 information:
  State                               Offline
```

show chassis pic fpc-slot pic-slot (FPC Offline)

```
user@host> show chassis pic fpc-slot 1 pic-slot 0
FPC 1 is not online
```

show chassis pic fpc-slot pic-slot (FPC Not Present)

```
user@host> show chassis pic fpc-slot 4 pic-slot 0
FPC slot 4 is empty
```

show chassis pic fpc-slot pic-slot (PIC Not Present)

```
user@host> show chassis pic fpc-slot 5 pic-slot 2
FPC 5, PIC 2 is empty
```

show chassis routing-engine

Syntax show chassis routing-engine
<bios | *slot*>

Syntax (J-EX Series Switch) show chassis routing-engine
<*slot*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display the status of the Routing Engine.

Options none—Display information about one or more Routing Engines.

bios—(Optional) Display the basic input/output system (BIOS) firmware version.

slot—(Systems with multiple Routing Engines) (Optional) Display information for an individual Routing Engine. Replace ***slot*** with 0 or 1.

Required Privilege Level view

Related Documentation
- **request chassis routing-engine master on page 210**

Output Fields Table 93 on page 653 lists the output fields for the **show chassis routing-engine** command. Output fields are listed in the approximate order in which they appear.

Table 93: show chassis routing-engine Output Fields

Field Name	Field Description
Slot	(Systems with multiple Routing Engines) Slot number.
Current state	(Systems with multiple Routing Engines) Current state of the Routing Engine: **Master**, **Backup**, or **Disabled**.
Election priority	(Systems with multiple Routing Engines) Election priority for the Routing Engine: **Master** or **Backup**.
Temperature	Temperature of the air flowing past the Routing Engine.
DRAM	Total DRAM available to the Routing Engine's processor.
Memory utilization	Percentage of Routing Engine memory being used.

Table 93: show chassis routing-engine Output Fields *(continued)*

Field Name	Field Description
CPU utilization	Information about the Routing Engine's CPU utilization: • **User**—Percentage of CPU time being used by user processes. • **Background**—Percentage of CPU time being used by background processes. • **Kernel**—Percentage of CPU time being used by kernel processes. • **Interrupt**—Percentage of CPU time being used by interrupts. • **Idle**—Percentage of CPU time that is idle.
Model	Routing Engine model number.
Serial ID	(Systems with multiple Routing Engines) Identification number of the Routing Engine in this slot.
Start time	Time at which the Routing Engine started running.
Uptime	How long the Routing Engine has been running.
Last reboot reason	Reason for last reboot, including: • **power cycle/failure**—Reboot due to the switching off of the power button behind the Routing Engine, not the power button on the chassis. • **watchdog**—Reboot due to a hardware watchdog. • **power-button hard power off**—Reboot due to pressing of the power button. • **misc hardware reason**—Reboot due to miscellaneous hardware reasons. • **thermal shutdown**—Reboot due to the router reaching a critical temperature point at which it is unsafe to continue operations. • **hard disk failure**—Reboot due to a hard disk failure. • **reset from debugger**—Reboot due to reset from the debugger. • **chassis control reset**—Reboot due to a chassis control reset. • **bios auto recovery reset**—Reboot due to a BIOS auto-recovery reset. • **could not be determined**—Reboot due to an undetermined reason. • **Router rebooted after a normal shutdown**—Reboot due to a normal shutdown.
Load averages	Routing Engine load averages for the last 1, 5, and 15 minutes.

show chassis temperature-thresholds

Syntax show chassis temperature-thresholds

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display chassis temperature threshold settings, in degrees Celsius.

Required Privilege Level view

List of Sample Output **show chassis temperature-thresholds on page 655**

Output Fields Table 94 on page 655 lists the output fields for the **show chassis temperature-thresholds** command. Output fields are listed in the approximate order in which they appear.

Table 94: show chassis temperature-thresholds Output Fields

Field name	Field Description
Item	Chassis component. If per FRU per slot thresholds are configured, the components about which information is displayed include the chassis, the Routing Engines, FPCs, and FEBs. If per FRU per slot thresholds are not configured, the components about which information is displayed include the chassis and the Routing Engines.
Fan speed	Temperature threshold settings, in degrees Celsius, for the fans to operate at normal and high speeds. • **Normal**—The fans operate at normal speed if the component is at or below this temperature and all the fans are present and functioning normally. • **High**—The fans operate at high speed if the component has exceeded this temperature or a fan has failed or is missing. An alarm is not triggered until the temperature exceeds the threshold settings for a yellow alarm or a red alarm.
Yellow alarm	Temperature threshold settings, in degrees Celsius, that trigger a yellow alarm. • **Normal**—The temperature that must be exceeded on the component to trigger a yellow alarm when the fans are running at full speed. • **Bad fan**—The temperature that must be exceeded on the component to trigger a yellow alarm when one or more fans have failed or are missing.
Red alarm	Temperature threshold settings, in degrees Celsius, that trigger a red alarm. • **Normal**—The temperature that must be exceeded on the component to trigger a red alarm when the fans are running at full speed. • **Bad fan**—The temperature that must be exceeded on the component to trigger a red alarm when one or more fans have failed or are missing.

show chassis temperature-thresholds

```
user@host> show chassis temperature-thresholds
                          Fan speed        Yellow alarm        Red alarm
Item                    Normal   High    Normal   Bad fan    Normal   Bad fan
Chassis default             48     54        65        55        75        65
Routing Engine 0            70     80        95        95       110       110
Routing Engine 1            70     80        95        95       110       110
```

FPC 0	55	60	75	65	90	80
FPC 1	55	60	75	65	90	80
FPC 2	55	60	75	65	90	80
FPC 3	55	60	75	65	90	80
FPC 4	55	60	75	65	90	80
FPC 5	55	60	75	65	90	80
FPC 6	55	60	75	65	90	80
FPC 7	55	60	75	65	90	80
FPC 8	55	60	75	65	90	80
FPC 9	55	60	75	65	90	80
FPC 10	55	60	75	65	90	80
FPC 11	55	60	75	65	90	80

show log

Syntax

show log
<*filename* | user <*username*>>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

List log files, display log file contents, or display information about users who have logged in to the router or switch.

Options

none—List all log files.

filename—(Optional) Display the log messages in the specified log file.

user <*username*>—(Optional) Display logging information about users who have recently logged in to the router or switch. If you include ***username***, display logging information about the specified user.

Required Privilege Level

trace

List of Sample Output

show log on page 657
show log filename on page 657
show log user on page 658

show log

```
user@host> show log
total 57518
-rw-r--r--  1 root  bin      211663 Oct  1 19:44 dcd
-rw-r--r--  1 root  bin      999947 Oct  1 19:41 dcd.0
-rw-r--r--  1 root  bin      999994 Oct  1 17:48 dcd.1
-rw-r--r--  1 root  bin      238815 Oct  1 19:44 rpd
-rw-r--r--  1 root  bin     1049098 Oct  1 18:00 rpd.0
-rw-r--r--  1 root  bin     1061095 Oct  1 12:13 rpd.1
-rw-r--r--  1 root  bin     1052026 Oct  1 06:08 rpd.2
-rw-r--r--  1 root  bin     1056309 Sep 30 18:21 rpd.3
-rw-r--r--  1 root  bin     1056371 Sep 30 14:36 rpd.4
-rw-r--r--  1 root  bin     1056301 Sep 30 10:50 rpd.5
-rw-r--r--  1 root  bin     1056350 Sep 30 07:04 rpd.6
-rw-r--r--  1 root  bin     1048876 Sep 30 03:21 rpd.7
-rw-rw-r--  1 root  bin       19656 Oct  1 19:37 wtmp
```

show log filename

```
user@host> show log rpd
Oct  1 18:00:18 trace_on: Tracing to ?/var/log/rpd? started
Oct  1 18:00:18 EVENT <MTU> ds-5/2/0.0 index 24 <Broadcast PointToPoint Multicast
Oct  1 18:00:18
Oct  1 18:00:19 KRT recv len 56 V9 seq 148 op add Type route/if af 2 addr
13.13.13.21 nhop type local nhop 13.13.13.21
Oct  1 18:00:19 KRT recv len 56 V9 seq 149 op add Type route/if af 2 addr
13.13.13.22 nhop type unicast nhop 13.13.13.22
Oct  1 18:00:19 KRT recv len 48 V9 seq 150 op add Type ifaddr index 24 devindex
43
Oct  1 18:00:19 KRT recv len 144 V9 seq 151 op chnge Type ifdev devindex 44
Oct  1 18:00:19 KRT recv len 144 V9 seq 152 op chnge Type ifdev devindex 45
Oct  1 18:00:19 KRT recv len 144 V9 seq 153 op chnge Type ifdev devindex 46
```

```
Oct  1 18:00:19 KRT recv len 1272 V9 seq 154 op chnge Type ifdev devindex 47
...
```

show log user

```
user@host> show log user
darius   mg2546                   Thu Oct  1 19:37   still logged in
darius   mg2529                   Thu Oct  1 19:08 - 19:36  (00:28)
darius   mg2518                   Thu Oct  1 18:53 - 18:58  (00:04)
root     mg1575                   Wed Sep 30 18:39 - 18:41  (00:02)
root     ttyp2    jun.site.per    Wed Sep 30 18:39 - 18:41  (00:02)
alex     ttyp1    192.168.1.2      Wed Sep 30 01:03 - 01:22  (00:19)
```

show pfe next-hop

Syntax
show pfe next-hop
<interface *interface-name*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display Packet Forwarding Engine next-hop information.

Options
none—Display all Packet Forwarding Engine next-hop information.

interface *interface-name*—(Optional) Display the Packet Forwarding Engine next-hop interface.

Required Privilege Level
admin

List of Sample Output
show pfe next-hop on page 660

show pfe next-hop

```
user@host> show pfe next-hop
Nexthop Info:
   ID      Type      Interface   Protocol         Encap  Next Hop Addr    MTU
-----  --------  -------------  ---------  ------------  ---------------  -----
    4     Mcast  -                   IPv4             -  0.0.0.0             0
    5     Bcast  -                   IPv4             -  -                   0
    7   Discard  -                   IPv4             -  -                   0
    8  MDiscard  -                   IPv4             -  -                   0
    9    Reject  -                   IPv4             -  -                   0
   13     Local  -                   IPv4             -  192.168.4.60        0
   14   Resolve  fxp0.0              IPv4   Unspecified  -                   0
   17     Local  -                   IPv4             -  127.0.0.1           0
   18   Unicast  fxp0.0              IPv4   Unspecified  192.168.4.254       0
   21     Local  -                   IPv4             -  11.1.0.1            0
   22   Unicast  at-0/1/0.0          IPv4      ATM SNAP  11.1.0.2          4482
 ...
```

show pfe route

Syntax

show pfe route
<<inet6 | ip | iso> <prefix *prefix*> | <table <*table-name*> <index *index*> <prefix *prefix*>>>
<mpls>
<summary>

Syntax (J-EX Series Switch)

show pfe route
<<inet6 | ip> <prefix *prefix*> | <table <*table-name*> <index *index*> <prefix *prefix*>>>
<mpls>
<summary>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display the routes in the Packet Forwarding Engine forwarding table. The Packet Forwarding Engine forwards packets between input and output interfaces.

NOTE: The Routing Engine maintains a master copy of the forwarding table. It copies the forwarding table to the Packet Forwarding Engine, which is the part of the router or switch responsible for forwarding packets. To display the routes in the Routing Engine forwarding table, use the **show route forwarding table** command. For more information, see the *Junos OS Routing Protocols and Policies Command Reference*.

Options

none—Display all Packet Forwarding Engine forwarding table information.

inet6—(Optional) Display Packet Forwarding Engine IPv6 routes.

ip—(Optional) Display Packet Forwarding Engine IPv4 routes.

iso —(Optional) Display ISO version routing tables.

mpls—(Optional) Display Packet Forwarding Engine Multiprotocol Label Switching (MPLS) information.

prefix *prefix*—(Optional) IPv4 or IPv6 prefix for which to show table entries.

summary—(Optional) Display summary of Packet Forwarding Engine information.

table <*table-name*> <index *index*> <prefix *prefix*>—(Optional) Display table information. Optionally, specify the table name, index, or prefix.

Required Privilege Level

admin

List of Sample Output

show pfe route ip on page 661
show pfe route iso on page 662

show pfe route ip

```
user@host> show pfe route ip

IPv4 Route Table 0, default.0, 0x0:
```

```
Destination                       NH IP Addr      Type     NH ID Interface
--------------------------------- --------------- -------- ----- ---------
default                                            Discard     8
127.0.0.1                         127.0.0.1          Local   256
172.16/12                         192.168.71.254   Unicast    68 fxp0.0
192.168.0/18                      192.168.71.254   Unicast    68 fxp0.0
192.168.40/22                     192.168.71.254   Unicast    68 fxp0.0
192.168.64/18                     192.168.71.254   Unicast    68 fxp0.0
192.168.64/21                                      Resolve    67 fxp0.0
192.168.71.249                    192.168.71.249     Local    66
192.168.220.0/30                                   Resolve   303 fe-0/0/0.0
192.168.220.0                     192.168.220.0    Receive   301 fe-0/0/0.0
224.0.0.1                                            Mcast     5
255.255.255.255                                      Bcast     6

...
```

show pfe route iso

```
user@host# show pfe route iso

CLNS Route Table 0, CLNP.0, 0x0:
Destination                                 Type     NH ID Interface
------------------------------------------- -------- ----- ---------
default                                       Reject    60
47.0005.80ff.f800.0000.0108.0001.0102.5508.2159/152     Local   514
49.0001.00a0.c96b.c491/72                      Local   536
```

show pfe statistics ip

Syntax show pfe statistics ip
<icmp | options>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display IPv4 Packet Forwarding Engine statistics.

Options none—Display all IPv4 Packet Forward Engine statistics.

icmp—(Optional) Display Packet Forwarding Engine IP ICMP statistics.

options—(Optional) Display Packet Forwarding Engine IP options statistics.

Required Privilege Level admin

List of Sample Output **show pfe statistics ip icmp on page 664**
show pfe statistics ip options on page 665

Output Fields Table 95 on page 663 lists the output fields for the **show pfe statistics ip** command. Output fields are listed in the approximate order in which they appear.

Table 95: show pfe statistics ip Output Fields

Field Name	Field Description
ICMP Statistics	ICMP statistics, including the following: • **requests**—Number of ICMP notifications sent to the PFE. If a throttler is configured, the number of notifications might not reflect all requests made. (See the **throttled icmps** field description.) • **network unreachables**—When route lookups fail, ICMP packets are sent to the source. These packets are ICMP TypeDestination Unreachable (3) and ICMP Code=Network Unreachable (0). • **ttl expired**—Number of notifications processed as a result of time-to-live (TTL) expiration packets. • **ttl captured**—Number of TTL expired packets sent by PFE interfaces to the Routing Engine. • **redirects**—Number of ICMP errors sent with Type=Redirect (5). • **mtu exceeded**—Number of ICMP errors sent with Type=Source Quench (4). • **icmp/option handoffs**—Number of packets that the PFE hardware requests the PFE software to process.

Table 95: show pfe statistics ip Output Fields *(continued)*

Field Name	Field Description
ICMP errors	ICMP errors, including the following: • **unknown unreachables**—Unknown code (greater than 16) found for an unknown unreachable type ICMP error. • **unsupported ICMP type**—Any ICMP type other than UNREACH, REDIRECT, TIME_EXCEED, and PARAM_PROB. • **unprocessed redirects**—When trying to find the neighbor to send redirects to, the PFE could not find the next-hop information. • **invalid ICMP type**—Any ICMP type other than UNREACH, REDIRECT, TIME_EXCEED, and PARAM_PROB. • **invalid protocol**—An incorrect protocol was detected by the ICMP processor. • **bad input interface if1**—The PFE software cannot map the interface index supplied by the chips to a proper data structure in the microkernel. • **throttled icmps**—Number of requests dropped because of rate limiting by the PFE. • **runts**—Number of packets for which the IP header length is less than the minimum length that is supported.
ICMP Discards	ICMP discard statistics, including the following: • **multicasts**—ICMP packets are not sent for link-layer multicast packets. These are counted as invalid source addresses (not a unicast address or all zeros). • **bad source addresses**—ICMP packets were received from an invalid source address (not a unicast address or all zeros). • **bad dest addresses**—ICMP packets were sent to an invalid destination address (not a unicast address or all zeros). • **IP fragments**—ICMP responses are sent only for the first fragments. The rest do not receive a response. This is the count for ICMP requests that receive no response. • **ICMP errors**—Number of ICMP error packets.

show pfe statistics ip icmp

```
user@host> show pfe statistics ip icmp
ICMP Statistics:
           0 requests
           0 network unreachables
           0 ttl expired
           0 ttl captured
           0 redirects
           0 mtu exceeded
           0 icmp/option handoffs
ICMP Errors:
           0 unknown unreachables
           0 unsupported ICMP type
           0 unprocessed redirects
           0 invalid ICMP type
           0 invalid protocol
           0 bad input interface
           0 throttled icmps
           0 runts
ICMP Discards:
           0 multicasts
           0 bad source addresses
           0 bad dest addresses
           0 IP fragments
```

```
            0 ICMP errors
```

show pfe statistics ip options

```
user@host> show pfe statistics ip options
IP Option Values:
  LSRR/SSRR forwarding enabled
IP Option Statistics:
            0 loose source routes
            0 strict source routes
            0 record routes
       889382 router alerts
            0 other options
IP Option Errors:
            0 runts
            2 bad versions
            0 runt header lengths
            0 giant header lengths
            0 null frames
            0 bad option lengths
            0 duplicate options
            0 bad option pointers
            0 source route frames dropped
          188 frames queued
         1126 frames dropped
```

show pfe statistics ip6

Syntax show pfe statistics ip6
<icmp>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display Packet Forwarding Engine IPv6 statistics.

Options none—Display all Packet Forwarding Engine IPv6 statistics.

icmp—(Optional) Display Packet Forwarding Engine IP ICMP statistics.

Required Privilege Level admin

List of Sample Output **show pfe statistics ip6 icmp on page 667**
show pfe statistics ip6 lcc on page 668

Output Fields Table 96 on page 666 lists the output fields for the **show pfe statistics ip6** command. Output fields are listed in the approximate order in which they appear.

Table 96: show pfe statistics ip6 Output Fields

Field Name	Field Description
ICMP6 Statistics	ICMP6 statistics, including the following: • **requests**—Number of ICMP notifications sent to the PFE. If a throttler is configured, the number of notifications might not reflect all requests made. (See the **throttled icmps** field description.) • **network unreachables**—When route lookups fail, ICMP packets are sent to the source. These packets are ICMP Type= Destination Unreachable (3) and ICMP Code= Network Unreachable (0). • **ttl expired**—Number of notifications processed as a result of time-to-live (TTL) expiration packets. • **ttl captured**—Number of TTL expired packets sent by PFE interfaces to the Routing Engine. • **redirects**—Number of ICMP errors sent with Type=Redirect (5). • **mtu exceeded**—Number of ICMP errors sent with Type=Source Quench (4). • **icmp/option handoffs**—Number of packets that the PFE hardware requests the PFE software to process.

Table 96: show pfe statistics ip6 Output Fields *(continued)*

Field Name	Field Description
ICMP6 errors	ICMP6 errors, including the following: • **unknown unreachables**—Unknown code (greater than 16) found for an unknown unreachable type ICMP error. • **unsupported ICMP type**—Any ICMP type other than UNREACH, REDIRECT, TIME_EXCEED, and PARAM_PROB. • **unprocessed redirects**—When trying to find the neighbor to send redirects to, the PFE could not find the next-hop information. • **invalid ICMP type**—Any ICMP type other than UNREACH, REDIRECT, TIME_EXCEED, and PARAM_PROB. • **invalid protocol**—An incorrect protocol was detected by the ICMP processor. • **bad input interface if1**—The PFE software cannot map the interface index supplied by the chips to a proper data structure in the microkernel. • **throttled icmps**—Number of requests dropped because of rate limiting by the PFE. • **runts**—Number of packets for which the IP header length is less than the minimum length that is supported.
ICMP6 Discards	ICMP6 discard statistics, including the following: • **multicasts**—ICMP packets are not sent for link-layer multicast packets. These are counted as invalid source addresses (not a unicast address or all zeros). • **bad source addresses**—ICMP packets were received from an invalid source address (not a unicast address or all zeros). • **bad dest addresses**—ICMP packets were sent to an invalid destination address (not a unicast address or all zeros). • **IP fragments**—ICMP responses are sent only for the first fragments. The rest do not receive a response. This is the count for ICMP requests that receive no response. • **ICMP errors**—Number of ICMP error packets.

show pfe statistics ip6 icmp

```
user@host> show pfe statistics ip6 icmp
ICMP6 Statistics:
           0 requests
           0 network unreachables
           0 ttl expired
           0 ttl captured
           0 redirects
           0 mtu exceeded
           0 icmp/option handoffs
ICMP6 Errors:
           0 unknown unreachables
           0 unsupported ICMP type
           0 unprocessed redirects
           0 invalid ICMP type
           0 invalid protocol
           0 bad input interface
           0 throttled icmps
           0 runts
ICMP6 Discards:
           0 multicasts
           0 bad source addresses
           0 bad dest addresses
```

```
                0 IP fragments
                0 ICMP errors
```

show pfe statistics ip6 lcc

```
user@host> show pfe statistics ip6 lcc 0 fpc 0
sfc0-re0:
--------------------------------------------------------------------------------
ICMP Statistics:
                0 requests
                0 network unreachables
                0 ttl expired
                0 ttl captured
                0 redirects
                0 mtu exceeded
                0 icmp/option handoffs

ICMP Errors:
                0 unknown unreachables
                0 unsupported ICMP type
                0 unprocessed redirects
                0 invalid ICMP type
                0 invalid protocol
                0 bad input interface
                0 throttled icmps
                0 runts

ICMP Discards:
                0 multicasts
                0 bad source addresses
                0 bad dest addresses
                0 IP fragments
                0 ICMP errors
```

show pfe terse

Syntax show pfe terse

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display Packet Forwarding Engine status information.

Options none—Display brief information about the Packet Forwarding Engine.

Required Privilege Level admin

show system alarms

Syntax show system alarms

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display active system alarms.

Options This command has no options.

Additional Information System alarms are preset. They include a **configuration** alarm that appears when no rescue configuration alarm is set and a **license** alarm that appears when a software feature is configured and no valid license is configured for the feature. For more information about system alarms, see the *Junos OS System Basics Configuration Guide*.

Required Privilege Level admin

List of Sample Output **show system alarms on page 670**

show system alarms

```
user@host> show system alarms
2 alarms currently active
Alarm time                    Class     Description
2005-02-24 17:29:34 UTC       Minor     IPsec VPN tunneling usage requires a
license
2005-02-24 17:29:34 UTC       Minor     Rescue configuration is not sent
```

show system audit

Syntax

show system audit
<root-only>

Syntax (J-EX Series Switch)

show system audit
<all-members>
<local>
<member *member-id*>
<root-only>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display the state and checksum values for file systems.

Options

none—Display the state and checksum values for all file systems.

all-members—(J-EX4200 switches only) (Optional) Display file system MD5 hash and permissions information on all members of the Virtual Chassis configuration.

local—(J-EX4200 switches only) (Optional) Display file system MD5 hash and permissions information on the local Virtual Chassis member.

member *member-id*—(J-EX4200 switches only) (Optional) Display file system MD5 hash and permissions information on the specified member of the Virtual Chassis configuration. Replace ***member-id*** with a value from 0 through 9.

root-only—(Optional) Check only the root (/) file system.

Additional Information

To redirect the output to a file, issue the following command:

ssh *router-name* 'show system audit root-only' > *output-file*

If you save the output of the **show system audit root-only** command to a file, you can compare it to subsequent output from the command to determine whether anything has changed.

Required Privilege Level

admin

List of Sample Output

show system audit root-only on page 671

show system audit root-only

```
user@host> show system audit root-only
#          user: root
#       machine: my-host
#          tree: /
date: Fri Feb 11 21:21:46 2000

# .
/set type=file uid=0 gid=0 mode=0755 nlink=1
.               type=dir nlink=23 size=1024 time=950252640.0
  .cshrc      uid=3 gid=7 mode=0644 size=177 time=939182975.0 \
               md5digest=f414e06fea6bd646244b98e13d6e6226
  .kernel.jkernel.backup \
```

```
              mode=0744 size=1934552 time=944688902.0 \
              md5digest=2c343cf0bd9fea8f04f78604feed7aa4
.profile    uid=3 gid=7 mode=0644 nlink=2 size=173 time=939182975.0 \
              md5digest=55a1e3c6c67789c9d3a1cce1ea39f670
  COPYRIGHT   uid=3 gid=7 mode=0444 size=3425 time=939182975.0 \
              md5digest=7df8bc77dcee71382ea73eb0ec6a9243
  boot.config mode=0644 size=3 time=945902618.0 \
              md5digest=93d722493ed38477338a1405d7dcbb40
  boot.help   uid=3 gid=7 mode=0444 size=411 time=939182876.0 \
              md5digest=9b7126385734bcae753f4179ab59d8e5
  compat      type=link mode=0777 size=11 time=915149058.0 \
              link=/usr/compat
  kernel      mode=0444 size=1947607 time=950230892.0 \
              md5digest=1a2a8aff2fec678a918ba0d6bf063980
  kernel.avr  uid=1112 size=1947642 time=950252597.0 \
              md5digest=82e1637682d58ec28964dfee7fccb62e
  kernel.config \
              mode=0644 size=0 time=915149058.0 \
              md5digest=d41d8cd98f00b204e9800998ecf8427e
  sys         type=link mode=0777 size=11 time=915149029.0 \
              link=usr/src/sys
```

show system buffers

Syntax show system buffers

Syntax (J-EX Series Switch)
show system buffers
<all-members>
<local>
<member *member-id*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display information about the buffer pool that the Routing Engine uses for local traffic. Local traffic is the routing and management traffic that is exchanged between the Routing Engine and the Packet Forwarding Engine within the router or switch, as well as the routing and management traffic from IP (that is, from OSPF, BGP, SNMP, ping operations, and so on).

Options none—Show all buffer statistics.

all-members—(J-EX4200 switches only) (Optional) Show buffer statistics for on all members of the Virtual Chassis configuration.

local—(J-EX4200 switches only) (Optional) Show buffer statistics for the local Virtual Chassis member.

member *member-id*—(J-EX4200 switches only) (Optional) Show buffer statistics for the specified member of the Virtual Chassis configuration. Replace ***member-id*** with a value from 0 through 9.

Additional Information A special type of memory buffer called a *cluster* is 2 KB in size. For more information, see *The Design and Implementation of the 4.4BSD Operation System* by McKusic, Bostic, Karels, and Quarterman.

Required Privilege Level view

List of Sample Output **show system buffers on page 674**

Output Fields Table 97 on page 674 describes the output fields for the **show system buffers** command. Output fields are listed in the approximate order in which they appear.

Table 97: show system buffers Output Fields

Field Name	Field Description
mbufs in use	Memory buffers (mbufs) are 128-byte buffers that are used for various purposes inside the kernel. Each memory buffer has a type, and the output itemizes the amount allocated for each type. Types with no memory buffers allocated are not displayed.
mbufs allocated to packet headers	Number of memory buffers currently holding packet headers
mbufs allocated to control blocks	Number of memory buffers currently holding state for sockets.
mbufs allocated to send data	Number of memory buffers currently holding socket send data.
mbufs allocated to pfe refill data	Number of memory buffers currently holding Packet Forwarding Engine refill data.
mbufs allocated to fxp data	Number of memory buffers currently holding fxp data.
mbufs allocated to socket names and addresses	Number of memory buffers currently holding addresses for sockets.
mbuf clusters in use	Allocation statistics for mbuf clusters.
allocated to network	Total amount of memory in use by the networking and interprocess communication (IPC) code.
requests for memory denied	Number of times a memory allocation request within the IPC and networking code failed.
requests for memory delayed	Number of times a memory allocation request within the IPC and networking code was postponed.
calls to protoocl drain routines	Number of times a memory allocation request within the IPC and networking code triggered a memory reclamation attempt.

show system buffers

```
user@host> show system buffers
        853 mbufs in use:
        2 mbufs allocated to packet headers
        37 mbufs allocated to protocol control blocks
        28 mbufs allocated to socket names and addresses
        2 mbufs allocated to socket send data
        400 mbufs allocated to pfe refill data
        384 mbufs allocated to fxp data
784/944 mbuf clusters in use
1994 Kbytes allocated to network (83% in use)
0 requests for memory denied
0 requests for memory delayed
0 calls to protocol drain routines
```

show system connections

Syntax

show system connections
<extensive>
<all-chassis | all-lcc | lcc *number* | scc>
<inet | inet6>
<show-routing-instances>

Syntax (J-EX Series Switch)

show system connections
<extensive>
<all-members>
<inet | inet6>
<local>
<member *member-id*>
<show-routing-instances>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display information about the active IP sockets on the Routing Engine. Use this command to verify which servers are active on a system and what connections are currently in progress.

Options

none—Display information about all active IP sockets on the Routing Engine.

extensive—(Optional) Display exhaustive system process information, which, for TCP connections, includes the TCP control block. This option is useful for debugging TCP connections.

all-members—(J-EX4200 switches only) (Optional) Display system connection activity for all members of the Virtual Chassis configuration.

inet | inet6—(Optional) Display IPv4 connections or IPv6 connections, respectively.

local—(J-EX4200 switches only) (Optional) Display system connection activity for the local Virtual Chassis member.

member *member-id*—(J-EX4200 switches only) (Optional) Display system connection activity for the specified member of the Virtual Chassis configuration. Replace ***member-id*** with a value from 0 through 9.

show-routing-instances—(Optional) Display routing instances.

Required Privilege Level

view

List of Sample Output

show system connections on page 676
show system connections extensive on page 677
show system connections show-routing-instances on page 677

Output Fields

Table 98 on page 676 describes the output fields for the **show system connections** command. Output fields are listed in the approximate order in which they appear.

Table 98: show system connections Output Fields

Field Name	Field Description
Proto	Protocol of the socket: **IP**, **TCP**, or **UDP** for IPv4 or IPv6.
Recv-Q	Number of input packets received by the protocol and waiting to be processed by the application.
Send-Q	Number of output packets sent by the application and waiting to be processed by the protocol.
Local Address	Local address and port of the socket, separated by a period. An asterisk (*) indicates that the bound address is the wildcard address. Server sockets typically have the wildcard address and a well-known port bound to them.
Foreign Address	Foreign address and port of the socket, separated by a period. An asterisk (*) indicates that the address or port is a wildcard.
Routing Instance (Displayed only when the **show-routing-instance** option is used.	Routing instances associated with active IP sockets on the Routing Engine.
(state)	For TCP, the protocol state of the socket.

show system connections

```
user@host> show system connections
Active Internet connections (including servers)
Proto Recv-Q Send-Q  Local Address          Foreign Address        (state)
tcp        0      2  192.168.4.16.513       208.197.169.254.894    ESTABLISHED
tcp        0      0  192.168.4.16.513       208.197.169.195.945    ESTABLISHED
tcp        0      0  *.23                   *.*                    LISTEN
tcp        0      0  *.22                   *.*                    LISTEN
tcp        0      0  *.513                  *.*                    LISTEN
tcp00 *.514      *.*                    LISTEN
tcp 0 0*.21      *.*                    LISTEN
tcp00 *.79     *.*                    LISTEN
tcp 00 *.1023     *.*                    LISTEN
tcp 00 *.111     *.*                    LISTEN
udp00192.168.4.16.1634     208.197.169.249.2049
udp00192.168.4.16.1627     208.197.169.254.2049
udp00192.168.4.16.1371     208.197.169.195.2049
udp00*.*      *.*
udp00*.9999     *.*
udp00 *.161     *.*
udp00192.168.4.16.1039     192.168.4.16.1023
udp00192.168.4.16.1038     192.168.4.16.1023
udp 00 192.168.4.16.1037     192.168.4.16.1023
udp00192.168.4.16.1036     192.168.4.16.1023
udp00*.1022     *.*
udp00*.1023     *.*
udp00*.111     *.*
udp00*.*      *.*
```

show system connections extensive

```
user@host> show system connections extensive
Active Internet connections (including servers)
Proto Recv-Q Send-Q  Local Address          Foreign Address        (state)
tcp        0      2  192.168.4.16.513       208.197.169.254.894    ESTABLISHED
       iss: 3972677059      sndup: 3972693435       sndcc:         10
    snduna: 3972693435     sndnxt: 3972693437      sndwnd:      17376
    sndmax: 3972693437    sndcwnd:      65535 sndssthresh: 1073725440
       irs:  484187869      rcvup:  484188060       rcvcc:       98357
    rcvnxt:  484188070     rcvadv:  484205446      rcvwnd:      17376
       rtt:          1       srtt:          7        rttv:          5
rxtcur:        120   rxtshift:          0       rtseq: 1103707591
rttmin:          2   duration:       5011         mss:       1448
     flags: REQ_SCALE RCVD_SCALE REQ_TSTMP RCVD_TSTMP RCVD_CC [0x41e0]
tcp        0      0  192.168.4.16.513       208.197.169.195.945    ESTABLISHED
       iss: 1057609890      sndup: 1057790796       sndcc:          2
    snduna: 1057790810     sndnxt: 1057790810      sndwnd:      17376
    sndmax: 1057790810    sndcwnd:      39096 sndssthresh: 1073725440
       irs: 3551947312      rcvup: 3551947422       rcvcc:          0
    rcvnxt: 3551947422     rcvadv: 3551964798      rcvwnd:      17376
       rtt:          0       srtt:         17        rttv:         11
rxtcur:        300   rxtshift:          0       rtseq:          0
rttmin:          2   duration:     125814         mss:       1448
     flags: REQ_SCALE RCVD_SCALE REQ_TSTMP RCVD_TSTMP [0x1e0]
udp0      0192.168.4.16.1634208.197.169.249.2049
udp0      0192.168.4.16.1627208.197.169.254.2049
udp0      0192.168.4.16.1371208.197.169.195.2049
udp 0      0*.* *.*
udp0      0*.9999*.*
udp 0      0*.161*.*
udp0      0192.168.4.16.1039192.168.4.16.1023
udp0      0192.168.4.16.1038192.168.4.16.1023
udp0      0192.168.4.16.1037192.168.4.16.1023
udp0      0192.168.4.16.1036192.168.4.16.1023
udp0      0*.1022*.*
udp 0      0*.1023 *.*
udp0      0 *.111*.*
udp0      0*.**.*
```

show system connections show-routing-instances

```
user@host> show system connections show-routing-instances
Active Internet connections (including servers) (including routing-instances)
Proto Recv-Q Send-Q  Local Address         Foreign Address       Routing Instance
      (state)
tcp4       0      0  192.168.69.204.23      172.17.28.19.4267      default
          ESTABLISHED
tcp4       0      0  192.168.69.204.58540   10.209.7.138.23        default
          ESTABLISHED
tcp4       0      0  192.168.69.204.23      172.17.28.19.1098      default
          ESTABLISHED
tcp4       0      0  192.168.7.1.57668      192.168.9.1.179        default
          ESTABLISHED
tcp4       0      0  192.168.7.1.179        192.168.8.1.49209      default
          ESTABLISHED
tcp4       0      0  128.0.0.1.6234         128.0.3.17.1024
__juniper_private1__   ESTABLISHED
tcp4       0      0  128.0.0.4.9000         128.0.0.4.59103
__juniper_private1__   ESTABLISHED
tcp4       0      0  128.0.0.4.59103        128.0.0.4.9000
__juniper_private1__   ESTABLISHED
tcp4       0      0  *.32012                *.*
__juniper_private1__   LISTEN
tcp4       0      0  *.9000                 *.*
```

```
__juniper_private1__   LISTEN
tcp4       0      0  *.33007                *.*
__juniper_private2__   LISTEN
tcp46      0      0  *.179                  *.*                    default
           LISTEN
tcp4       0      0  *.179                  *.*                    default
           LISTEN
tcp4       0      0  *.6154                 *.*
__juniper_private1__   LISTEN
tcp4       0      0  *.6153                 *.*
__juniper_private1__   LISTEN
tcp4       0      0  *.7000                 *.*
__juniper_private1__   LISTEN
tcp4       0      0  *.6152                 *.*
__juniper_private1__   LISTEN
tcp4       0      0  *.6156                 *.*
__juniper_private1__   LISTEN
tcp4       0      0  *.33005                *.*
__juniper_private2__   LISTEN
tcp4       0      0  *.31343                *.*
__juniper_private1__   LISTEN
tcp4       0      0  *.31341                *.*
__juniper_private1__   LISTEN
tcp4       0      0  *.32003                *.*
__juniper_private2__   LISTEN
tcp4       0      0  *.666                  *.*
__juniper_private1__   LISTEN
tcp4       0      0  *.38                   *.*
__juniper_private1__   LISTEN
tcp4       0      0  *.3221                 *.*                    default
           LISTEN
```

show system core-dumps

Syntax

show system core-dumps
<brief | detail>
<*core-filename*>
<core-file-info>

Syntax (J-EX Series Switch)

show system core-dumps
<all-members>
<brief | detail>
<*core-filename*>
<core-file-info>
<local>
<member *member-id*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Show core files on all routers or switches running the Junos OS. You can use the **show system core-dumps** command to show a list of system core files created when the router or switch has failed. This command can be useful for diagnostic purposes. Each list item includes the file permissions, number of links, owner, group, size, modification date, and path/filename.

You can use the option ***core-filename*** and its options **core-file-info**, **brief**, and **detail** to display more information about the specified core-dump files.

Options

none—Display a list of all existing core-dump files.

all-members—(J-EX4200 switches only) (Optional) Display system core files on all members of the Virtual Chassis configuration.

brief—(Optional) View details of binary.

core-file-info—(Optional) Display the stack trace of a core file.

core-filename—(Optional) Name of a specific core file to display.

detail—(Optional) View stack trace with details of binary.

local—(J-EX4200 switches only) (Optional) Display system core files on the local Virtual Chassis member.

member *member-id*—(J-EX4200 switches only) (Optional) Display system core files on the specified member of the Virtual Chassis configuration. Replace ***member-id*** with a value from 0 through 9.

Required Privilege Level

view

List of Sample Output

show system core-dumps on page 680
show system core-dumps on page 680

Output Fields Table 99 on page 680 describes the output fields for the **show system core-dumps** command. Output fields are listed in the approximate order in which they appear.

Table 99: show system core-dumps Output Fields

Field Name	Field Description
Permissions	Read/write permissions for the file named.
Links	Number of links to the file.
Owner	Name of the file owner.
Group	Name of the group with file access.
File size	File size in bytes.
Modified	Last file modification date and time.
Path/filename	File path where the file resides and the filename.

show system core-dumps This example shows the command output if core files exist.

```
user@host> show system core-dumps
-rw------- 1 root wheel 268369920 Jun 18 17:59 /var/crash/vmcore.0
-rw-rw---- 1 root field 3371008 Jun 18 17:53 /var/tmp/rpd.core.0
-rw-r--r-- 1 root wheel 27775914 Jun 18 17:59 /var/crash/kernel.0
```

show system core-dumps This example shows the command output if core files do not exist.

```
user@host> show system core-dumps
/var/crash/*core*: No such file or directory
/var/tmp/*core*: No such file or directory
/var/crash/kernel.*: No such file or directory
```

show system directory-usage

Syntax

show system directory-usage
<depth *number*>
<*path*>

Syntax (J-EX Series Switch)

show system directory-usage
<all-members>
<depth *number*>
<local>
<member *member-id*>
<*path*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display directory usage information.

Options

none—Display all directory usage information.

all-members—(J-EX4200 switches only) (Optional) Display directory information for all members of the Virtual Chassis configuration.

depth *number*—(Optional) Depth of the directory to traverse. This option is useful when you want to limit the output shown for a large file system.

local—(J-EX4200 switches only) (Optional) Display directory information for the local Virtual Chassis member.

member *member-id*—(J-EX4200 switches only) (Optional) Display directory information for the specified member of the Virtual Chassis configuration. Replace ***member-id*** with a value from 0 through 9.

path—(Optional) Path or root directory to traverse.

Required Privilege Level

view

Output Fields

Table 100 on page 681 describes the output fields for the **show system directory-usage** command. Output fields are listed in the approximate order in which they appear.

Table 100: show system directory-usage Output Fields

Field Name	Field Description
bytes	Number of bytes used by files in a directory.
directory-name	Name of the directory.

show system processes

Syntax

show system processes
<brief | detail | extensive | summary>
<health (pid *process-identifer* | process-name *process-name*)>
<providers>
<resource-limits (brief | detail) *process-name*>
<wide>

Syntax (J-EX Series Switch)

show system processes
<all-members>
<brief | detail | extensive | summary>
<health (pid *process-identifer* | process-name *process-name*)>
<local>
<member *member-id*>
<providers>
<resource-limits (brief | detail) *process-name*>
<wide>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display information about software processes that are running on the router or switch and that have controlling terminals.

Options

none—Display standard information about system processes.

all-members—(J-EX4200 switches only) (Optional) Display standard system process information for all members of the Virtual Chassis configuration.

brief | detail | extensive | summary—(Optional) Display the specified level of detail.

health (pid *process-identifer* | process-name *process-name*)—(Optional) Display process health information.

local—(J-EX4200 switches only) (Optional) Display standard system process information for the local Virtual Chassis member.

member *member-id*—(J-EX4200 switches only) (Optional) Display standard system process information for the specified member of the Virtual Chassis configuration. Replace ***member-id*** with a value from 0 through 9.

providers—(Optional) Display provider processes.

resource-limits (brief | detail) *process-name*—(Optional) Display process resource limits.

wide—(Optional) Display process information that might be wider than 80 columns.

Required Privilege Level

view

List of Sample Output

show system processes on page 685
show system processes brief on page 685
show system processes detail on page 685

show system processes extensive on page 686
show system processes summary on page 687

Output Fields Table 101 on page 683 describes the output fields for the **show system processes** command. Output fields are listed in the approximate order in which they appear.

Table 101: show system processes Output Fields

Field Name	Field Description	Level of Output
last PID	Last process identifier assigned to the process.	**brief extensive summary**
load averages	Three load averages followed by the current time.	**brief extensive summary**
processes	Number of existing processes and the number of processes in each state (**sleeping, running, starting, zombies**, and **stopped**).	**brief extensive summary**
Mem	Information about physical and virtual memory allocation.	**brief extensive summary**
Swap	Information about physical and virtual memory allocation.	**brief extensive summary**
PID	Process identifier.	**detail extensive summary**
TT	Control terminal name.	none **detail**

Table 101: show system processes Output Fields *(continued)*

Field Name	Field Description	Level of Output
STAT	Symbolic process state. The state is given by a sequence of letters. The first letter indicates the run state of the process: • **D**—In disk or other short-term, uninterruptible wait • **I**—Idle (sleeping longer than about 20 seconds) • **R**—Runnable • **S**—Sleeping for less than 20 seconds • **T**—Stopped • **Z**—Dead (zombie) • **+** —The process is in the foreground process group of its control terminal. • **<** —The process has raised CPU scheduling priority. • **>** —The process has specified a soft limit on memory requirements and is currently exceeding that limit; such a process is not swapped. • **A**—The process requested random page replacement. • **E**—The process is trying to exit. • **L**—The process has pages locked in core. • **N**—The process has reduced CPU scheduling priority. • **S**—The process requested first-in, first-out (FIFO) page replacement. • **s**—The process is a session leader. • **V**—The process is temporarily suspended. • **W**—The process is swapped out. • **X**—The process is being traced or debugged.	none **detail**
UID	User identifier.	**detail**
USERNAME	Process owner.	**extensive summary**
PPID	Parent process identifier.	**detail**
CPU	(D)—Short-term CPU usage. (E and S)—Raw (unweighted) CPU usage. The value of this field is used to sort the processes in the output.	**detail extensive summary**
RSS	Resident set size.	**detail**
WCHAN	Symbolic name of the wait channel.	**detail**
STARTED	Local time when the process started running.	**detail**
PRI	Current priority of the process. A lower number indicates a higher priority.	**detail extensive summary**
NI or NICE	UNIX "niceness" value. A lower number indicates a higher priority.	**detail extensive summary**
SIZE	Total size of the process (text, data, and stack), in kilobytes.	**extensive summary**

Table 101: show system processes Output Fields *(continued)*

Field Name	Field Description	Level of Output
RES	Current amount of resident memory, in kilobytes.	extensive summary
STATE	Current state of the process (for example, **sleep**, **wait**, **run**, **idle**, **zombie**, or **stop**).	extensive summary
TIME	(S)—Number of system and user CPU seconds that the process has used. (None, D, and E)—Total amount of time that the command has been running.	detail extensive summary
WCPU	Weighted CPU usage.	extensive summary
COMMAND	Command that is currently running.	detail extensive summary

show system processes

```
user@host> show system processes
PID  TT  STAT      TIME COMMAND
    0  ??  DLs    0:00.70  (swapper)
    1  ??  Is     0:00.35 /sbin/init --
    2  ??  DL     0:00.00  (pagedaemon)
    3  ??  DL     0:00.00  (vmdaemon)
    4  ??  DL     0:42.37  (update)
    5  ??  DL     0:00.00  (if_jnx)
   80  ??  Ss     0:14.66 syslogd -s
   96  ??  Is     0:00.01 portmap
  128  ??  Is     0:02.70 cron
  173  ??  Is     0:02.24 /usr/local/sbin/sshd (sshd1)
  189  ??  S      0:03.80 /sbin/watchdog -t180
  190  ??  I      0:00.03 /usr/sbin/tnetd -N
  191  ??  S      2:24.76 /sbin/ifd -N
  192  ??  S<     0:55.44 /usr/sbin/xntpd -N
  195  ??  S      0:53.11 /usr/sbin/snmpd -N
  196  ??  S      1:15.73 /usr/sbin/mib2d -N
  198  ??  I      0:00.75 /usr/sbin/inetd -N
 2677  ??  I      0:00.01 /usr/sbin/mgd -N
 2712  ??  Ss     0:00.24 rlogind
 2735  ??  R      0:00.00 /bin/ps -ax
 1985  p0- S      0:07.41 ./rpd -N
 2713  p0  Is     0:00.24 -tcsh (tcsh)
 2726  p0  S+     0:00.07 cli
```

show system processes brief

```
user@host> show system processes brief
last pid:   543;  load averages:  0.00,  0.00,  0.00    18:29:47
37 processes:  1 running, 36 sleeping

Mem: 25M Active, 3976K Inact, 19M Wired, 8346K Buf, 202M Free
Swap: 528M Total, 64K Used, 528M Free
```

show system processes detail

```
user@host> show system processes detail
PID   UID  PPID CPU PRI NI  RSS WCHAN  STARTED   TT  STAT      TIME COMMAND
 3151  1049  3129   2  28  0  672 -       1:13PM   p0  R+     0:00.00 ps -ax -r
    1     0     0   0  10  0  376 wait    1:51PM   ??  Is     0:00.29 /sbin/ini
    2     0     0   0 -18  0   12 psleep  1:51PM   ??  DL     0:00.00  (pagedae
    3     0     0   0  28  0   12 psleep  1:51PM   ??  DL     0:00.00  (vmdaemo
    4     0     0   0  28  0   12 update  1:51PM   ??  DL     0:07.15  (update)
```

```
    5     0     0   0   2  0   12 pfesel  1:51PM   ??  IL     0:02.90  (if_pfe)
   27     0     1   0  10  0 17936 mfsidl  1:51PM  ??  Is     0:00.46 mfs /dev/
   81     0     1   0   2  0  496 select  1:52PM   ??  Ss     0:31.21 syslogd -
  119     1     1   0   2  0  492 select  1:52PM   ??  Is     0:00.00 portmap
  134     0     1   0   2  0  580 select  1:52PM   ??  S      0:02.95 amd -p -a
  151     0     1   0  18  0  532 pause   1:52PM   ??  Is     0:00.34 cron
  183     0     1   0   2  0  420 select  1:52PM   ??  Ss     0:00.07 /usr/loca
  206     0     1   0  18  0   72 pause   1:52PM   ??  S      0:00.51 /sbin/wat
  207     0     1   0   2  0  520 select  1:52PM   ??  I      0:00.16 /usr/sbin
  208     0     1   0   2  0  536 select  1:52PM   ??  S      0:08.21 /sbin/dcd
  210     0     1 255   2 -12  740 select  1:52PM  ??  S<     0:05.83 /usr/sbin
  211     0     1   0   2  0  376 select  1:52PM   ??  S      0:00.03 /usr/sbin
  215     0     1   0   2  0  548 select  1:52PM   ??  I      0:00.50 /usr/sbin
  219     0     1   0   3  0  540 ttyin   1:52PM   v0  Is+    0:00.02 /usr/libe
  220     0     1   0   3  0  540 ttyin   1:52PM   v1  Is+    0:00.01 /usr/libe
  221     0     1   0   3  0  540 ttyin   1:52PM   v2  Is+    0:00.01 /usr/libe
  222     0     1   0   3  0  540 ttyin   1:52PM   v3  Is+    0:00.01 /usr/libe
  735     0     1   0   2  0  468 select  2:47PM   ??  S      0:19.14 /usr/sbin
  736     0     1   0   2  0  212 select  2:47PM   ??  S      0:14.13 /usr/sbin
 1380     0     1   0   3  0  888 ttyin   7:32PM   d0  Is+    0:00.46 bash
 3019     0   207   0   2  0  636 select 10:49AM   ??  Ss     0:02.93 tnp.chass
 3122     0  1380   0   2  0 1764 select 12:33PM   d0  S      0:00.77 ./rpd -N
 3128     0   215   0   2  0  580 select 12:45PM   ??  Ss     0:00.12 rlogind
 3129  1049  3128   0  18  0  944 pause  12:45PM   p0  Ss     0:00.14 -tcsh (tc
    0     0     0   0 -18  0    0 sched   1:51PM   ??  DLs    0:00.10  (swapper
```

show system processes extensive

```
user@host> show system processes extensive
last pid:   544;  load averages:  0.00,  0.00,  0.00    18:30:33
37 processes:  1 running, 36 sleeping

Mem: 25M Active, 3968K Inact, 19M Wired, 8346K Buf, 202M Free
Swap: 528M Total, 64K Used, 528M Free
  PID USERNAME PRI NICE SIZE    RES STATE    TIME   WCPU    CPU COMMAND
  544 root     30   0   604K   768K RUN      0:00  0.00%  0.00% top
    3 root     28   0     0K    12K psleep   0:00  0.00%  0.00% vmdaemon
    4 root     28   0     0K    12K update   0:03  0.00%  0.00% update
  528 aviva    18   0   660K   948K pause    0:00  0.00%  0.00% tcsh
  204 root     18   0   300K   544K pause    0:00  0.00%  0.00% csh
  131 root     18   0   332K   532K pause    0:00  0.00%  0.00% cron
  186 root     18   0   196K    68K pause    0:00  0.00%  0.00% watchdog
   27 root     10   0   512M 16288K mfsidl   0:00  0.00%  0.00% mount_mfs
    1 root     10   0   620K   344K wait     0:00  0.00%  0.00% init
  304 root      3   0   884K   900K ttyin    0:00  0.00%  0.00% bash
  200 root      3   0   180K   540K ttyin    0:00  0.00%  0.00% getty
  203 root      3   0   180K   540K ttyin    0:00  0.00%  0.00% getty
  202 root      3   0   180K   540K ttyin    0:00  0.00%  0.00% getty
  201 root      3   0   180K   540K ttyin    0:00  0.00%  0.00% getty
  194 root      2   0  2248K  1640K select   0:11  0.00%  0.00% rpd
  205 root      2   0   964K   800K select   0:12  0.00%  0.00% tnp.chassisd
  189 root      2 -12   352K   740K select   0:03  0.00%  0.00% xntpd
  114 root      2   0   296K   612K select   0:00  0.00%  0.00% amd
  188 root      2   0   780K   600K select   0:00  0.00%  0.00% dcd
  527 root      2   0   176K   580K select   0:00  0.00%  0.00% rlogind
  195 root      2   0   212K   552K select   0:00  0.00%  0.00% inetd
  187 root      2   0   192K   532K select   0:00  0.00%  0.00% tnetd
   83 root      2   0   188K   520K select   0:00  0.00%  0.00% syslogd
  538 root      2   0  1324K   516K select   0:00  0.00%  0.00% mgd
   99 daemon    2   0   176K   492K select   0:00  0.00%  0.00% portmap
  163 root      2   0   572K   420K select   0:00  0.00%  0.00% nsrexecd
  192 root      2   0   560K   400K select   0:10  0.00%  0.00% snmpd
  191 root      2   0  1284K   376K select   0:00  0.00%  0.00% mgd
```

```
  537 aviva     2   0   636K   364K select   0:00  0.00%  0.00% cli
  193 root      2   0   312K   204K select   0:07  0.00%  0.00% mib2d
    5 root      2   0     0K    12K pfesel   0:00  0.00%  0.00% if_pfe
    2 root    -18   0     0K    12K psleep   0:00  0.00%  0.00% pagedaemon
    0 root    -18   0     0K     0K sched    0:00  0.00%  0.00% swapper
```

show system processes summary

```
user@host> show system processes summary
last pid:   543;  load averages:  0.00,  0.00,  0.00    18:29:47
37 processes:  1 running, 36 sleeping

Mem: 25M Active, 3976K Inact, 19M Wired, 8346K Buf, 202M Free
Swap: 528M Total, 64K Used, 528M Free

  PID USERNAME PRI NICE SIZE    RES STATE    TIME   WCPU    CPU COMMAND
  527 root      2   0   176K   580K select   0:00  0.04%  0.04% rlogind
  543 root     30   0   604K   768K RUN      0:00  0.00%  0.00% top
```

PART 11

Virtual Chassis

CHAPTER 43

Virtual Chassis—Overview, Components, and Configurations

Virtual Chassis Overview

The Dell PowerConnect J-Series J-EX4200 Ethernet Switch is the basis for the *Virtual Chassis* flexible, scaling switch solution. You can connect individual J-EX4200 switches together to form one unit and manage the unit as a single chassis, called a Virtual Chassis. Up to ten J-EX4200 switches can be interconnected, providing up to a total of 480 access ports. The available bandwidth increases as you include more members within the Virtual Chassis configuration. See "Understanding the High-Speed Interconnection of the Virtual Chassis Members" on page 702.

This topic describes:

- Global Management of Member Switches in a Virtual Chassis on page 693
- High Availability Through Redundant Routing Engines on page 693
- Adaptability as an Access Switch or Distribution Switch on page 693

Basic Configuration of a Virtual Chassis with Master and Backup Switches

To take advantage of the Virtual Chassis configuration's higher bandwidth capacity and software redundancy features, you need to interconnect at least two J-EX4200 switches in a Virtual Chassis configuration. You can start with a default configuration, composed of two J-EX4200 *member switches* interconnected through the dedicated 64-Gbps *Virtual Chassis ports (VCPs)* on their rear panels. These ports do not have to be configured. They are operational as soon as the member switches are powered on. See "Example: Configuring a Virtual Chassis with a Master and Backup in a Single Wiring Closet" on page 717 for additional information.

Expanding Configurations—Within a Single Wiring Closet and Across Wiring Closets

As your needs grow, you can easily expand the Virtual Chassis configuration to include more member switches. Within a single wiring closet, simply add member switches by cabling together the dedicated VCPs. For more information about expanding Virtual Chassis configurations within a single wiring closet, see "Example: Expanding a Virtual Chassis Configuration in a Single Wiring Closet" on page 722 and "Example: Setting Up a Multimember Virtual Chassis Access Switch with a Default Configuration" on page 727.

You can also expand a Virtual Chassis configuration beyond a single wiring closet. Interconnect switches located in multiple wiring closets or in multiple data center racks by installing the optional SFP or SFP+ uplink modules and connecting the uplink module ports or by connecting the 1-gigabit network interfaces in a J-EX4200-24F switch. The small form-factor pluggable (SFP) uplink module provides four ports for 1-gigabit transceivers. The SFP+ uplink module provides two ports for 10-gigabit SFP+ transceivers or four ports for 1-gigabit SFP transceivers. To use SFP and SFP+ uplink module ports or J-EX4200-24F network interfaces for interconnecting member switches, you must first explicitly configure them as *Virtual Chassis ports (VCPs)*. This procedure includes configuring these ports of a standalone J-EX4200 switch as VCPs prior to interconnecting the new member switch with the existing Virtual Chassis configuration. See "Example: Configuring a Virtual Chassis Interconnected Across Multiple Wiring Closets" on page 733 for detailed information.

When you are creating a Virtual Chassis configuration with multiple members, you might want to deterministically control the role and member ID assigned to each member switch. You can do this by creating a preprovisioned configuration. See "Example: Configuring a Virtual Chassis Using a Preprovisioned Configuration File" on page 752 for more information.

You can add switches to a preprovisioned configuration by using the autoprovisioning feature to automatically configure the uplink module ports as VCPs on the switches being added. See "Adding a New Switch to an Existing Virtual Chassis Configuration (CLI Procedure)" on page 786 for detailed information.

Global Management of Member Switches in a Virtual Chassis

The interconnected member switches in a Virtual Chassis configuration operate as a single network entity. You run EZSetup only once to specify the identification parameters for the master, and these parameters implicitly apply to all members of the Virtual Chassis configuration. You can view the Virtual Chassis configuration as a single device in the J-Web user interface and apply various device management functions to all members of the Virtual Chassis configuration.

The serial console port and dedicated out-of-band management port that are on the rear panel of the individual switches have global virtual counterparts when the switches are interconnected in a Virtual Chassis configuration. A *virtual console* allows you to connect to the master by connecting a terminal directly to the console port of any member switch. A *virtual management Ethernet (VME)* interface allows you to remotely manage the Virtual Chassis configuration by connecting to the out-of-band management port of any member switch through a single IP address. See “Understanding Global Management of a Virtual Chassis Configuration” on page 699.

High Availability Through Redundant Routing Engines

A Virtual Chassis configuration has a master and a backup, each of which has a Routing Engine. These redundant Routing Engines handle all routing protocol processes and control the Virtual Chassis configuration. See “High Availability Features for J-EX Series Switches Overview” on page 18 for further information on redundant Routing Engines and additional high availability features.

Adaptability as an Access Switch or Distribution Switch

A Virtual Chassis configuration supports a variety of user environments, because it can be composed of different models of J-EX4200 switches, with either 24 or 48 access ports, and with these having either full (24 or 48 ports) or partial (8 ports) Power over Ethernet (PoE) port capabilities. You can select different switch models to support various functions. For example, you might set up one Virtual Chassis access switch configuration composed of the full PoE models to support users sitting in cubicles equipped with PCs and VoIP phones. You could set up another Virtual Chassis configuration with partial PoE models to support the company's internal servers and configure one more Virtual Chassis configuration with partial PoE models to support the company's external servers. Alternatively, the Virtual Chassis configuration can be used as a distribution switch. For this type of deployment, you might select the J-EX4200-24F model to connect the distribution switch to multiple access switches located in different buildings on the campus.

Related Documentation

- Understanding Virtual Chassis Components on page 694
- Understanding How the Master in a Virtual Chassis Configuration Is Elected on page 698
- Understanding Virtual Chassis J-EX4200 Switch Version Compatibility on page 705
- Understanding Virtual Chassis Configurations and Link Aggregation on page 702
- Understanding Virtual Chassis Configuration on page 704
- J-EX4200 Switch Models on page 26

Understanding Virtual Chassis Components

A Virtual Chassis configuration allows you to interconnect two to ten J-EX4200 Ethernet Switches and run them as a single network entity. While it is true that you need at least two interconnected switches to take advantage of Virtual Chassis features, it is also true that any individual J-EX4200 switch has some Virtual Chassis components.

This topic covers:

Virtual Chassis Ports (VCPs)

There are two dedicated Virtual Chassis ports (VCPs) on the rear panel of the J-EX4200 switch that are used exclusively to interconnect J-EX4200 switches in a Virtual Chassis configuration. The interfaces for these dedicated ports are operational by default when the ports are properly cabled. For an example of two J-EX4200 switches interconnected with their dedicated VCPs, see “Example: Configuring a Virtual Chassis with a Master and Backup in a Single Wiring Closet” on page 717. In addition, you can interconnect the switch with another J-EX4200 switch across a wider distance by installing an optional SFP or SFP+ uplink module in a J-EX4200 switch or by using the network interfaces in a J-EX4200-24F switch. To do this using uplink module ports, you need to install one uplink module in at least one J-EX4200 switch at each end of the link. You must set the uplink module ports or the J-EX4200-24F network interfaces to function as VCPs in order for the interconnected switches to be recognized as members of the same Virtual Chassis configuration. This procedure includes setting the uplink module ports or J-EX4200-24F network ports of a standalone J-EX4200 switch as VCPs prior to interconnecting the new member switch with the existing Virtual Chassis configuration. For an example of J-EX4200 switches interconnected with the uplink ports functioning as VCPs, see “Example: Configuring a Virtual Chassis Interconnected Across Multiple Wiring Closets” on page 733.

You can display the status of both the dedicated VCP interfaces and the uplink ports configured as VCP interfaces with the **show virtual-chassis vc-port** command.

Master Role

The member that functions in the master role:

- Manages the member switches.
- Runs the Junos OS for J-EX Series Switches in a master role.

- Runs the chassis management processes and control protocols.
- Represents all the member switches interconnected within the Virtual Chassis configuration. (The hostname and other properties that you assign to this switch during setup apply to all members of the Virtual Chassis configuration.)

When a J-EX4200 switch is powered on as a standalone switch, it is considered the master member. In a multimember Virtual Chassis configuration, one member functions as the master and a second member functions as the backup:

- In a preprovisioned configuration, one of the two members assigned as **routing-engine** functions as the master member. The selection of which member assigned as **routing-engine** functions as master and which as backup is determined by the software based on the master election algorithm. See “Understanding How the Master in a Virtual Chassis Configuration Is Elected” on page 698.
- In a configuration that is not preprovisioned, the selection of the master and backup is determined by the mastership priority value and secondary factors in the master election algorithm.

Backup Role

The member that functions in the backup role:

- Maintains a state of readiness to take over the master role if the master fails.
- Runs the Junos OS for J-EX Series switches in a backup role.
- Synchronizes with the master in terms of protocol states, forwarding tables, and so forth, so that it is prepared to preserve routing information and maintain network connectivity without disruption in case the master is unavailable.

You must have at least two member switches in a Virtual Chassis configuration in order to have a backup member.

- In a preprovisioned configuration, one of the two members assigned as **routing-engine** functions in the backup role. The selection of which member assigned as **routing-engine** functions as master and which as backup is determined by the software based on the master election algorithm. See “Understanding How the Master in a Virtual Chassis Configuration Is Elected” on page 698.
- In a configuration that is not preprovisioned, the selection of the master and backup is determined by the mastership priority value and secondary factors in the master election algorithm.

Linecard Role

A member that functions in the linecard role:

- Runs only a subset of the Junos OS for J-EX Series switches.
- Does not run the chassis control protocols.
- Can detect certain error conditions (such as an unplugged cable) on any interfaces that have been configured on it through the master.

A Virtual Chassis configuration must have at least three members in order to include a linecard member.

- In a preprovisioned configuration, you can explicitly configure a member with the **role** of linecard, which makes it ineligible for functioning as a master or backup.
- In a configuration that is not preprovisioned, the members that are not selected as master or backup function as linecard members of the Virtual Chassis configuration. The selection of the master and backup is determined by the mastership priority value and secondary factors in the master election algorithm.

Member Switch and Member ID

Each physically discrete J-EX4200 switch is a potential member of a Virtual Chassis configuration. When a J-EX4200 switch is powered on, it receives a member ID that is displayed on the front-panel LCD. If the switch is powered on as a standalone switch, its member ID is always **0**. When the switch is interconnected with other J-EX4200 switches in a Virtual Chassis configuration, its member ID (**0** through **9**) is assigned by the master based on various factors, such as the order in which the switch was added to the Virtual Chassis configuration. As each switch is added and powered on, it receives the next available (unused) member ID.

If the Virtual Chassis configuration previously included a member switch and that member was physically disconnected or removed from the Virtual Chassis configuration, its member ID is not available for assignment as part of the standard sequential assignment by the master. For example, you might have a Virtual Chassis configuration composed of member 0, member 2, and member 3, because member 1 was removed. When you add another member switch and power it on, the master assigns it as member 4. However, you can use the **request virtual-chassis renumber** command to explicitly change the member ID of the new member switch to use member ID 1.

The member ID distinguishes the member switches from one another. You use the member ID:

- To assign a mastership priority value to a member switch
- To configure interfaces for a member switch
- To apply some operational commands to a member switch
- To display status or characteristics of a member switch

Mastership Priority

In a configuration that is not preprovisioned, you can designate the role (master, backup, or linecard) that a member switch performs within the Virtual Chassis configuration by configuring its mastership priority (from **1** to **255**). The mastership priority value is the factor with the highest precedence for selecting the master of the Virtual Chassis configuration.

The default value for mastership priority is **128**. When a J-EX4200 switch is powered on, it receives the default mastership priority value. Because it is the only member of the Virtual Chassis configuration, it is also the master. When you interconnect a standalone

switch to an existing Virtual Chassis configuration (which implicitly includes its own master), we recommend that you explicitly configure the mastership priority of the members that you want to function as the master and backup.

We recommend that you specify the same mastership priority value for both the master and backup members.

NOTE: Configuring the same mastership priority value for both the master and backup helps to ensure a smooth transition from master to backup in case the master becomes unavailable. It prevents the old master from preempting control from the backup in situations where the backup has taken control of the Virtual Chassis configuration due to the original master being unavailable.

We also recommend that you configure the highest possible mastership priority value (255) for those two members, because that guarantees that these two members continue to function as the master and backup when other members are added to the Virtual Chassis configuration. Any other members of the Virtual Chassis configuration (members with lower mastership priority) are considered linecard members.

In a preprovisioned configuration, the mastership priority value is assigned by the software, based on the specified role.

Virtual Chassis Identifier (VCID)

All members of a Virtual Chassis configuration share one Virtual Chassis identifier (VCID). This identifier is derived from internal parameters. When you are monitoring a Virtual Chassis configuration, the VCID is displayed in the user interface.

Related Documentation

- Virtual Chassis Overview on page 691
- Example: Configuring a Virtual Chassis with a Master and Backup in a Single Wiring Closet on page 717
- Example: Configuring a Virtual Chassis Interconnected Across Multiple Wiring Closets on page 733
- Example: Configuring a Virtual Chassis Using a Preprovisioned Configuration File on page 752
- Setting an Uplink Module Port as a Virtual Chassis Port (CLI Procedure) on page 792
- Command Forwarding Usage with a Virtual Chassis Configuration on page 803

Understanding How the Master in a Virtual Chassis Configuration Is Elected

All switches that are interconnected in a Virtual Chassis configuration are member switches of that Virtual Chassis. Each Virtual Chassis configuration has one member that functions as the *master* and controls the Virtual Chassis configuration.

When a Virtual Chassis configuration boots, the Junos OS for J-EX Series Switches automatically runs a master election algorithm to determine which member switch takes the role of master.

The algorithm that the software uses to determine the master is as follows:

1. Choose the member with the highest user-configured mastership priority (255 is the highest possible value).
2. Choose the member that was master the last time the Virtual Chassis configuration booted.
3. Choose the member that has been included in the Virtual Chassis configuration for the longest period of time. (For this to be a deciding factor, there has to be a minimum time lapse of one minute between the power-ons of the individual interconnected member switches.)
4. Choose the member with the lowest MAC address.

The variations among switch models, such as whether the switch has 48 or 24 ports, do not impact the master election algorithm. To ensure that a specific member is elected as the master:

1. Power on only the switch that you want to configure as master of the Virtual Chassis configuration.
2. Configure the mastership priority of that member to have the highest possible value (255).
3. Continue to configure other members through the master member, as desired.
4. Power on the other members.

Related Documentation

- Virtual Chassis Overview on page 691
- Understanding Virtual Chassis Components on page 694
- Understanding Virtual Chassis Configuration on page 704

Understanding Software Upgrade in a Virtual Chassis Configuration

A Virtual Chassis configuration can be composed of multiple J-EX4200 Ethernet Switches and each member switch is running Junos OS packages. For ease of management, the Virtual Chassis configuration provides flexible methods to upgrade software releases.

A new software release can be upgraded to the entire Virtual Chassis configuration or to a particular member in the Virtual Chassis configuration through a CLI or J-Web command.

A user can add software packages to either a single member of the Virtual Chassis configuration or to all members of the Virtual Chassis configuration at the same time.

Related Documentation

- Virtual Chassis Overview on page 691
- Understanding Virtual Chassis Components on page 694
- Understanding Automatic Software Update on Virtual Chassis Member Switches on page 715
- Installing Software on a J-EX Series Switch with a Single Routing Engine (CLI Procedure) on page 70

Understanding Global Management of a Virtual Chassis Configuration

A Virtual Chassis configuration is composed of multiple J-EX4200 Ethernet Switches, so it has multiple console ports and multiple out-of-band management Ethernet ports located on the rear panels of the switches.

You can connect a PC or laptop directly to a console port of any member switch to set up and configure the Virtual Chassis. When you connect to the console port of any member switch, the console session is redirected to the master switch, as shown in Figure 7 on page 700.

Figure 7: Console Session Redirection

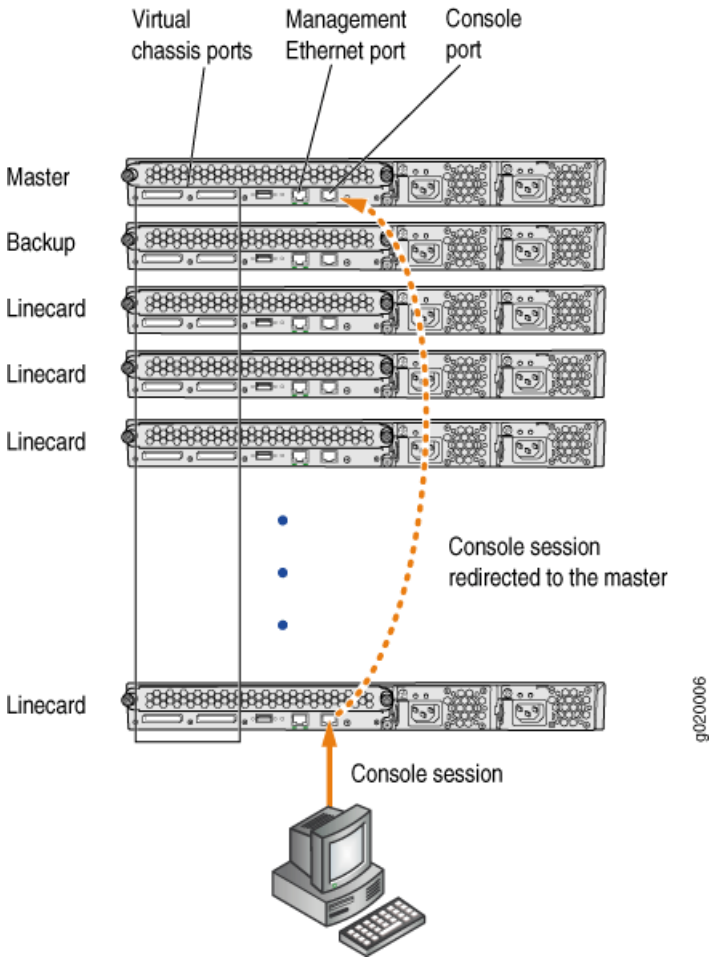

If the master becomes unavailable, the console session is disconnected from the old master and a new session is established with the newly elected master.

An out-of-band management Ethernet port is often referred to simply as a management Ethernet port. It uses a dedicated management channel for device maintenance and allows a system administrator to monitor and manage the switch by remote control.

The Virtual Chassis configuration can be managed remotely through SSH or Telnet using a global management interface called the virtual management Ethernet (VME) interface. VME is a logical interface representing any and all of the out-of-band management ports on the member switches. When you connect to the Virtual Chassis configuration using the VME IP address, the connection is redirected to the master member as shown in Figure 8 on page 701.

Figure 8: Management Ethernet Port Redirection to VME

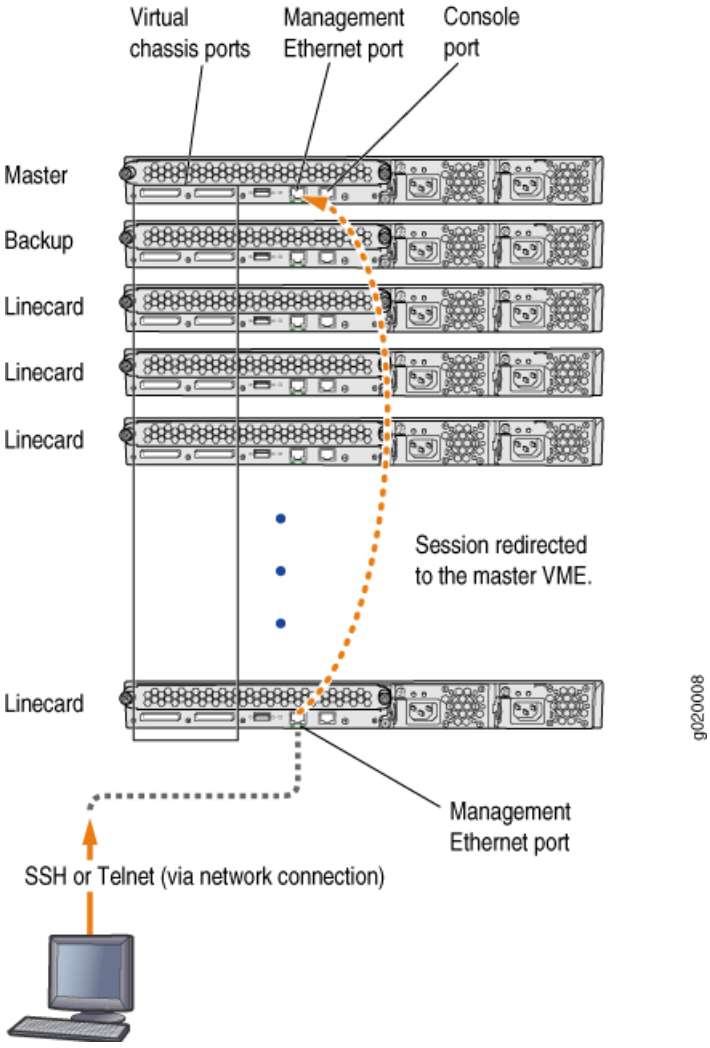

If the master management Ethernet link is unavailable, the session is redirected through the backup management Ethernet link. If there is no active management Ethernet link on the backup, the VME interface chooses a management Ethernet link on one of the linecard members, selecting the linecard member with the lowest member ID as its first choice.

You can configure an IP address for the VME global management interface at any time.

You can perform remote configuration and administration of all members of the Virtual Chassis configuration through the VME interface.

Related Documentation

- Understanding Virtual Chassis Components on page 694
- Example: Configuring a Virtual Chassis with a Master and Backup in a Single Wiring Closet on page 717

- Configuring the Virtual Management Ethernet Interface for Global Management of a Virtual Chassis (CLI Procedure) on page 797

Understanding Nonvolatile Storage in a Virtual Chassis Configuration

The J-EX4200 Ethernet Switch stores Junos OS system files in internal flash memory. In a Virtual Chassis configuration, both the master and the backup switch store the configuration information for all the member switches.

- Nonvolatile Memory Features on page 702

Nonvolatile Memory Features

The Junos OS for J-EX Series Switches optimizes the way the Virtual Chassis stores its configuration if a member switch or the Virtual Chassis configuration is shut down improperly:

- If the master is not available, the backup switch takes on the role of the master and its internal flash memory takes over as the alternate location for maintaining nonvolatile configuration memory.
- If a member switch is taken offline for repair, the master stores the configuration of the member switch.

Related Documentation

- Command Forwarding Usage with a Virtual Chassis Configuration on page 803
- Monitoring System Properties on page 550

Understanding the High-Speed Interconnection of the Virtual Chassis Members

Two high-speed Virtual Chassis ports (VCPs) on the rear panel of the Virtual Chassis member switches enable the members to be interconnected and operate as a single, powerful switch. Each VCP interface is 32 Gbps bidirectional. When VCP interfaces are used to form a ring topology, each segment provides 64 Gbps bidirectional bandwidth. Because the VCP links act as point-to-point links, multiple segments of the ring can be used simultaneously. This allows the Virtual Chassis configuration bandwidth to scale as you interconnect more members within the ring topology.

Related Documentation

- Understanding Virtual Chassis Components on page 694
- Virtual Chassis Cabling Configuration Examples for J-EX4200 Switches

Understanding Virtual Chassis Configurations and Link Aggregation

You can combine physical Ethernet ports belonging to different member switches of a Virtual Chassis configuration to form a logical point-to-point link, known as a *link aggregation group (LAG)* or *bundle*. A LAG provides more bandwidth than a single Ethernet link can provide. Additionally, link aggregation provides network redundancy by

load-balancing traffic across all available links. If one of the links fails, the system automatically load-balances traffic across all remaining links.

You can select up to four uplink module ports or SFP network ports on a J-EX4200-24F switch that have been configured as Virtual Chassis ports (VCPs) to form a LAG. When you set uplink module ports or SFP network ports on Virtual Chassis member switches as uplink VCPs, connect at least two of those uplink VCPs on one member to at least two uplink VCPs on another member, and configure those uplink VCPs to operate at the same link speed, the uplink VCPs automatically form a LAG and each LAG is assigned a positive-integer identifier called a *trunk ID*.

A LAG over uplink VCPs provides higher overall bandwidth for forwarding traffic between the member switches connected by the uplink VCPs, faster management communications, and greater redundancy of operations among the members than would be available without the LAG. All J-EX4200 Ethernet Switches have two dedicated VCPs. A LAG over uplink VCPs provides an additional Virtual Chassis link throughput of 20 Gbps for the J-EX4200-24T and J-EX4200-48T models and additional throughput of 28 Gbps for the J-EX4200-24F model. Up to eight Virtual Chassis LAGs can be created per member.

See “Setting an Uplink Module Port as a Virtual Chassis Port (CLI Procedure)” on page 792 for information about configuring uplink module ports and SFP network ports on J-EX4200-24F switches as uplink VCPs.

To verify that the LAG has been created, view the output of the command **show virtual-chassis vc-port**.

NOTE: The interfaces that are included within a bundle or LAG are sometimes referred to as *member interfaces*. Do not confuse this term with *member switches*, which refers to J-EX4200 switches that are interconnected as a Virtual Chassis. It is possible to create a LAG that is composed of member interfaces that are located in different member switches of a Virtual Chassis.

Related Documentation

- Virtual Chassis Overview on page 691
- Understanding Aggregated Ethernet Interfaces and LACP on page 867
- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 740
- Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 746
- Example: Configuring a Virtual Chassis Interconnected Across Multiple Wiring Closets on page 733
- Example: Configuring Link Aggregation Groups Using Uplink Virtual Chassis Ports on page 769

Understanding Virtual Chassis Configuration

You configure and manage almost all aspects of a Virtual Chassis configuration through the master of the Virtual Chassis. However, you can also configure Virtual Chassis parameters when a J-EX4200 Ethernet Switch is a standalone switch not interconnected with other members.

A J-EX4200 switch has some innate characteristics of a Virtual Chassis by default. A standalone J-EX4200 switch is assigned member ID **0** and is the master of itself. Therefore, you can edit its Virtual Chassis configuration. When the standalone switch is interconnected with an existing Virtual Chassis configuration, the Virtual Chassis configuration statements and any uplink Virtual Chassis port (VCP) settings that you previously specified on the standalone switch remain part of its configuration.

A switch is not recognized as a member of a Virtual Chassis until it is interconnected with the master or interconnected with an existing member of the Virtual Chassis. When a switch is located too far away to be interconnected through dedicated VCPs, you can specify an uplink module port or a J-EX4200-24F network interface as a VCP by using the **request virtual-chassis vc-port** command. You must issue the **request virtual-chassis vc-port** command on the switch you are adding to the Virtual Chassis as well as on the existing member switch that you will connect to the new member. Because the to-be-added switch is not yet a member, the master switch will not recognize that added switch unless the latter has an uplink VCP. A link aggregation group (LAG) will be formed automatically when the new switch is added to the configuration if more than one such link with the same speed is detected between uplink VCPs on the new member and an existing member. See “Understanding Virtual Chassis Configurations and Link Aggregation” on page 702.

When an uplink module port or a J-EX4200-24F network interface is set as a VCP, it cannot be used for any additional purpose. If you want to use the uplink module port or J-EX4200-24F network interface for another purpose, you can delete the VCP setting by using the **request virtual-chassis vc-port** command. You can execute this command directly on the member whose uplink VCP setting you want to delete or through the master of the Virtual Chassis configuration.

CAUTION: Deleting a VCP in a Virtual Chassis chain configuration can cause the Virtual Chassis configuration to split. For more information, see “Understanding Split and Merge in a Virtual Chassis Configuration” on page 712.

You can create a preprovisioned configuration. This type of configuration allows you to deterministically control the member ID and role assigned to a member switch by associating the switch with its serial number. For an example of a preprovisioned configuration, see “Example: Configuring a Virtual Chassis Using a Preprovisioned Configuration File” on page 752.

NOTE: If a J-EX4200 switch is interconnected with other switches in a Virtual Chassis configuration, each individual switch that is included as a member of the configuration is identified with a member ID. The member ID functions as an FPC slot number. When you are configuring interfaces for a Virtual Chassis configuration, you specify the appropriate member ID (0 through 9) as the *slot* element of the interface name.

The default factory settings for a Virtual Chassis configuration include FPC 0 as a member of the default VLAN because FPC 0 is configured as part of the ethernet-switching family. In order to include FPC 1 through FPC 9 in the default VLAN, add the ethernet-switching family to the configurations for those interfaces.

Related Documentation

- Understanding Virtual Chassis Components on page 694
- Understanding How the Master in a Virtual Chassis Configuration Is Elected on page 698
- Example: Configuring a Virtual Chassis Interconnected Across Multiple Wiring Closets on page 733
- Example: Configuring a Virtual Chassis with a Master and Backup in a Single Wiring Closet on page 717
- **request virtual-chassis vc-port on page 840**

Understanding Virtual Chassis J-EX4200 Switch Version Compatibility

For J-EX4200 Ethernet Switches to be interconnected as a Virtual Chassis configuration, the switches must be running the same software versions. The master checks the hardware version, the Junos OS version, and other component versions running in a switch that is physically interconnected to its Virtual Chassis port (VCP). Different hardware models can be members of the same Virtual Chassis configuration. However, the master will not assign a member ID to a switch that is running a different software version. A switch that is running a different version of software will not be allowed to join the Virtual Chassis configuration.

Related Documentation

- Understanding Virtual Chassis Components on page 694
- Understanding Software Upgrade in a Virtual Chassis Configuration on page 698
- Understanding Software Installation on J-EX Series Switches on page 61
- Installing Software on a J-EX Series Switch with a Single Routing Engine (CLI Procedure) on page 70
- Installing Software on J-EX Series Switches (J-Web Procedure) on page 75

Understanding Fast Failover in a Virtual Chassis Configuration

The Virtual Chassis fast failover feature is a hardware-assisted failover mechanism that automatically reroutes traffic and reduces traffic loss in the event of a link failure or switch failure. If a link between two members fails, traffic flow between those members must be rerouted quickly so that there is minimal traffic loss.

Fast failover is effective only for Virtual Chassis members configured in ring topologies using identical port types.

This topic describes the following:

Supported Topologies for Fast Failover

For fast failover to be effective, the Virtual Chassis members must be configured in a ring topology. The ring topology can be formed by using either dedicated Virtual Chassis ports (VCPs) or user-configured uplink module VCPs. Fast failover is supported only in a ring topology that uses identical port types, for example, either a topology that uses all dedicated VCPs or one that uses all uplink module VCPs. Fast failover is not supported in a ring topology that includes both dedicated VCPs and uplink module VCPs. Fast failover is supported, however, in a Virtual Chassis configuration that consists of multiple rings.

How Fast Failover Works

When fast failover is activated, each VCP is automatically configured with a backup port of the same type (dedicated VCP or SFP uplink VCP). If a VCP fails, its backup port is used to send traffic. These backup ports act as standby ports and are not meant for load-balancing purposes.

Fast Failover in a Ring Topology using Dedicated VCPs

When fast failover is activated in a ring topology that uses dedicated VCPs, each VCP is automatically configured with a backup port of the same type. If a VCP fails, its backup port is used to send traffic. Figure 9 on page 707 shows normal traffic flow in a ring topology using dedicated VCPs.

Figure 9: Normal Traffic Flow in a Ring Topology Using Dedicated VCPs

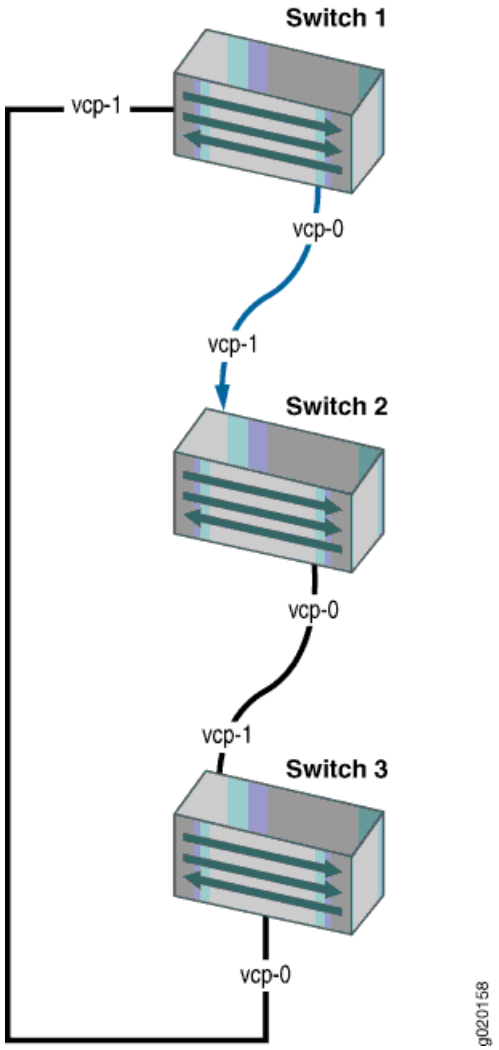

Figure 10 on page 708 shows traffic redirected by fast failover.

Figure 10: Traffic Redirected by Fast Failover After Dedicated VCP Link Failure

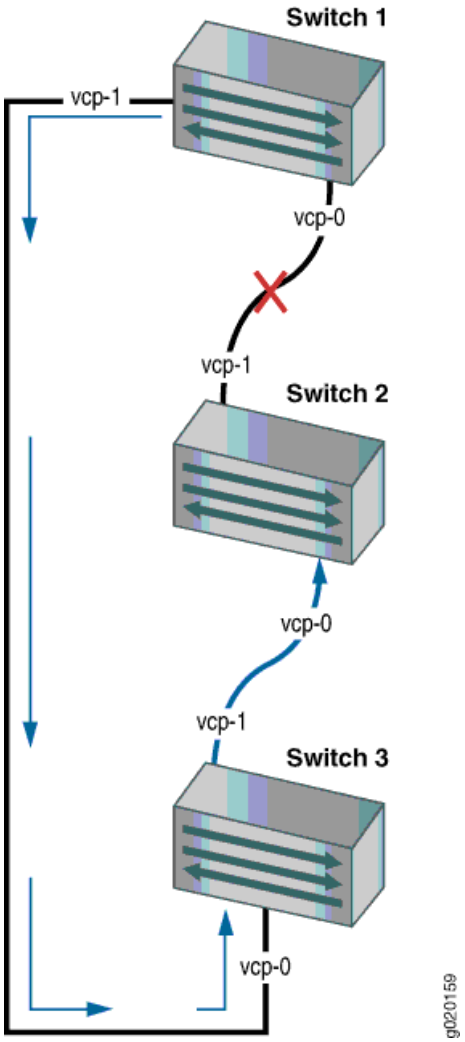

When the failed link is restored, the Virtual Chassis reconfigures the topology to the topology's original state.

Fast Failover in a Ring Topology Using Uplink Module VCPs

In a ring topology that uses uplink module VCPs, each uplink module VCP is automatically configured with a backup uplink module VCP. If an uplink module VCP fails, its backup port is used to send traffic. Figure 11 on page 709 shows normal traffic flow in a ring topology using SFP uplink module VCPs.

NOTE: In order to use SFP uplink module ports as VCPs, you must configure them to be VCPs using the request virtual-chassis vc-port command. Once configured, they will be converted into VCPs. For example xe-0/1/0 will become vcp-255/1/0 after you configure it to be a VCP.

Figure 11: Normal Traffic Flow in a Ring Topology Using SFP Uplink Module VCPs

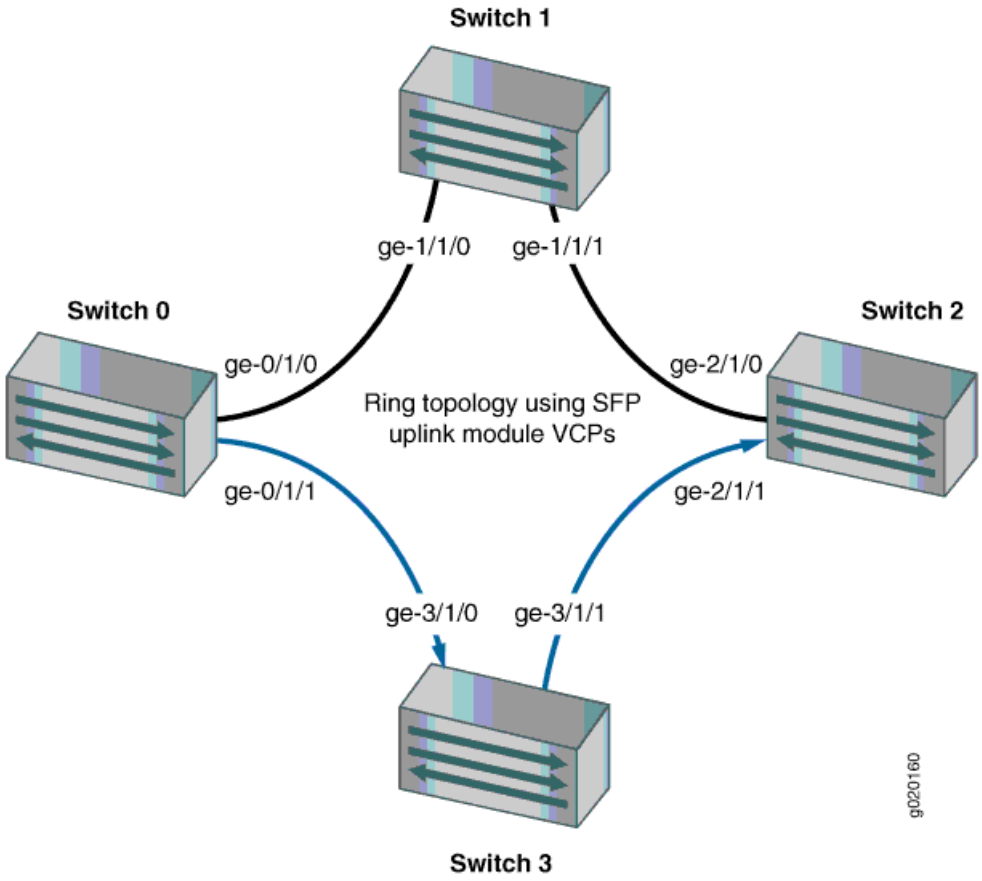

Figure 12 on page 710 shows traffic redirected by fast failover.

Figure 12: Traffic Redirected by Fast Failover After SFP Uplink Module VCP Link Failure

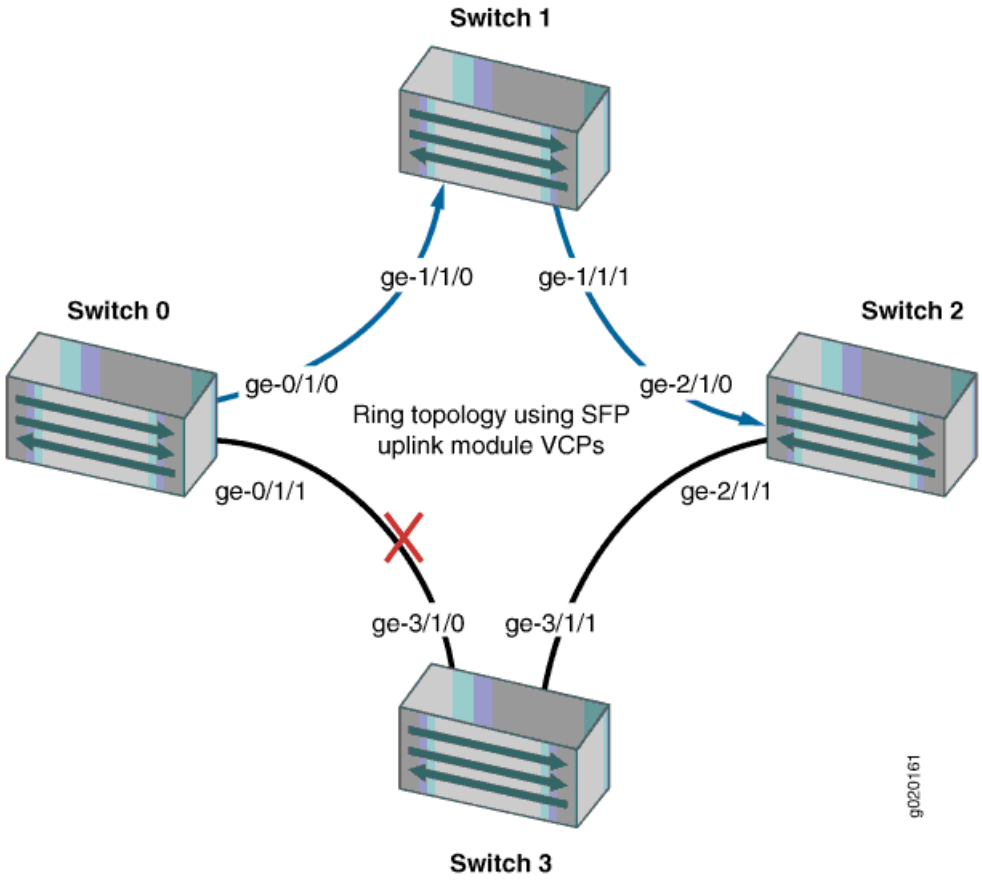

In a ring topology that uses SFP uplink module VCPs, there are four ports per module. Consecutive pair of ports are automatically configured as backup ports for each other. For example, if a Virtual Chassis member has an SFP uplink module installed, uplink module VCPs **ge-0/1/0** and **ge-0/1/1** are automatically configured as the backup port for the other port in the pair. Similarly, ports **ge-0/1/2** and **ge-0/1/3** are automatically configured as the backup port for the other port in the pair.

Fast Failover in a Virtual Chassis Configuration Using Multiple Ring Topologies

Fast failover is supported in a Virtual Chassis configuration with a multiple-ring topology, as shown in Figure 13 on page 711.

Figure 13: Traffic Redirected by Fast Failover After VCP Link Failures in a Topology with Multiple Rings

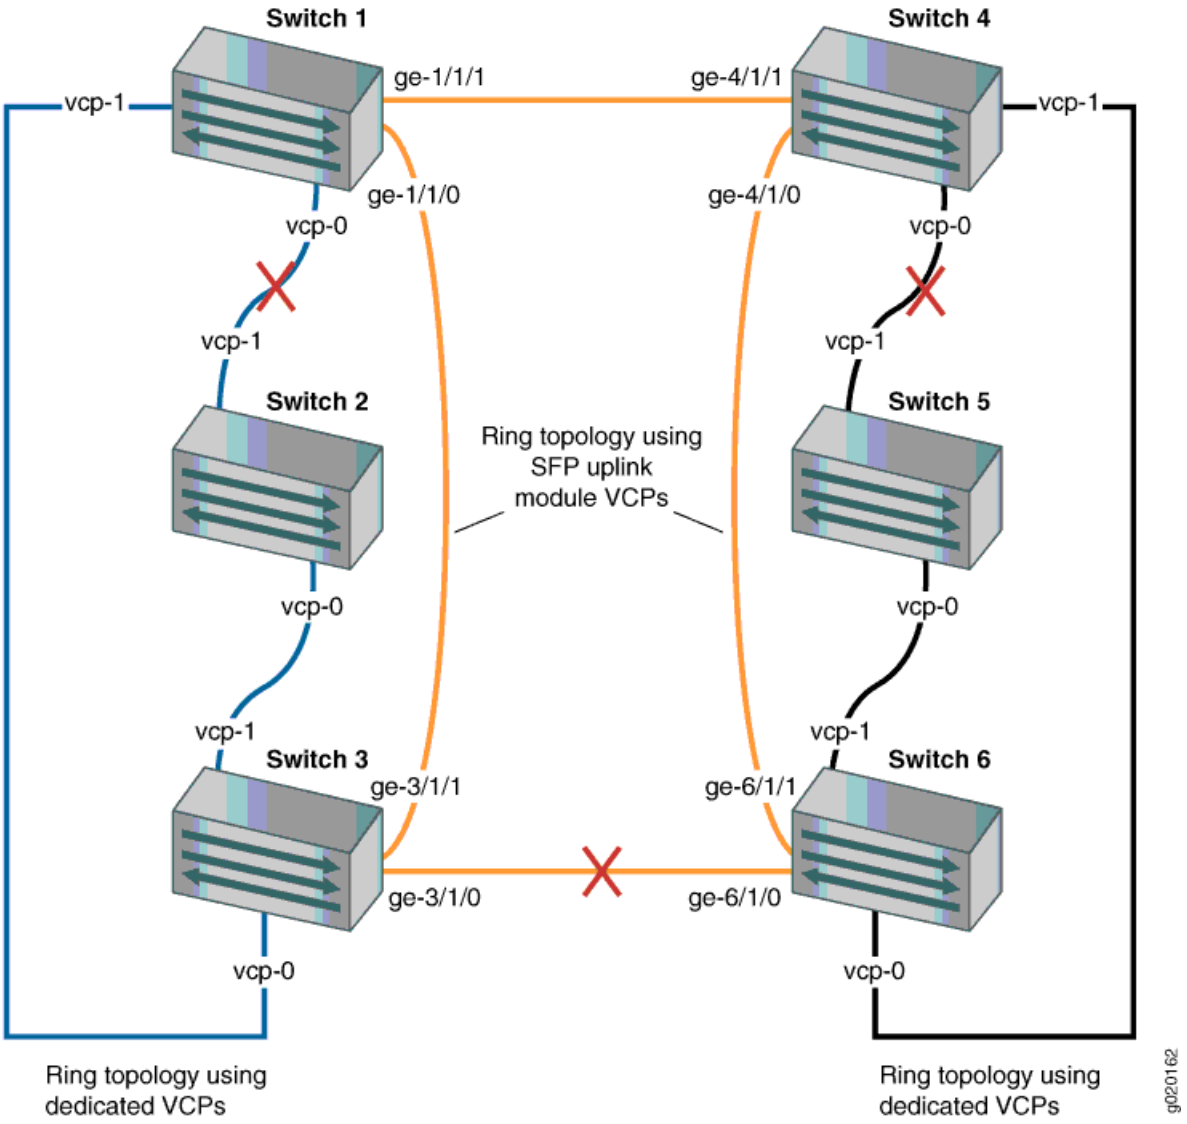

In this scenario, the Virtual Chassis configuration has three rings: two rings that use dedicated VCPs and one ring that uses SFP uplink module VCPs. Fast failover works independently on each ring. Each dedicated VCP in a ring is backed up by another dedicated VCP. Similarly, each SFP uplink module VCP is backed up by another SFP uplink module VCP. Fast failover does not support a ring topology consisting of a mix of dedicated VCPs and uplink module VCPs.

Effects of Topology Changes on a Fast Failover Configuration

Once the fast failover feature has been activated, topology changes to the Virtual Chassis configuration do not affect the fast failover configuration. In the event of a link or switch failure, fast failover functions normally.

Related Documentation

- Understanding Virtual Chassis Configuration on page 704
- Example: Configuring Fast Failover on Uplink Module VCPs to Reroute Traffic When a Virtual Chassis Member Switch or Intermember Link Fails on page 763
- Setting an Uplink Module Port as a Virtual Chassis Port (CLI Procedure) on page 792

Understanding Split and Merge in a Virtual Chassis Configuration

In a Virtual Chassis configuration, two or more J-EX4200 Ethernet Switches are connected together to form a unit that is managed as a single chassis. If there is a disruption to the Virtual Chassis configuration due to member switches failing or being removed from the configuration, the Virtual Chassis configuration splits into two separate Virtual Chassis. This situation could cause disruptions in the network if the two separate configurations share common resources, such as global IP addresses. The split and merge feature provides a method to prevent the separate Virtual Chassis configurations from adversely affecting the network and also allows the two parts to merge back into a single Virtual Chassis configuration.

NOTE: If a Virtual Chassis configuration splits into separate parts, we recommend that you resolve the problem that caused the Virtual Chassis configuration to split as soon as possible.

You can also use this feature to merge two active but separate Virtual Chassis that have not previously been part of the same configuration into one Virtual Chassis configuration.

NOTE: The split and merge feature is enabled by default on J-EX4200 switches. You can disable the split and merge feature by using the **set virtual-chassis no-split-detection** command.

This topic describes:

- What Happens When a Virtual Chassis Configuration Splits on page 712
- Merging Virtual Chassis Configurations on page 713

What Happens When a Virtual Chassis Configuration Splits

When a Virtual Chassis configuration splits into two separate Virtual Chassis configurations, the individual member switches detect this topology change and run the master election algorithm to select a new master for each of the two Virtual Chassis configurations. The new masters then determine whether their Virtual Chassis configuration remains active. One of the configurations remains active based on the following:

- It contains both the stable master and the stable backup (that is, the master and backup from the original Virtual Chassis configuration before the split).

- It contains the stable master and the configuration is greater than half the Virtual Chassis size.
- It contains the stable backup and is at least half the Virtual Chassis size.

Due to the rules given in the second and third list items, if the Virtual Chassis configuration splits into two equal parts and the stable master and stable backup are in different parts, then the part that contains the stable backup will become active.

NOTE: The number of members in the Virtual Chassis configuration includes all member switches connected to date minus the number whose Virtual Chassis member IDs have been recycled. Therefore, the size of the Virtual Chassis configuration increases when a new member switch is detected and decreases when a member switch's ID is recycled (that is, made available for reassignment).

These rules ensure that only one of the two separate Virtual Chassis configurations created by the split remains active. The member switches in the inactive Virtual Chassis configuration remain in a linecard role. For the inactive members to become active again, one of the following things must happen:

- The problem that caused the original Virtual Chassis configuration to split is resolved, allowing the two Virtual Chassis configurations to merge.
- You load the factory default configuration on the inactive members, which causes the inactive members to function as standalone switches or become part of a different Virtual Chassis configuration.

NOTE: When you remove a member switch from a Virtual Chassis configuration, you should recycle the member ID using the request virtual-chassis recycle command.

Merging Virtual Chassis Configurations

There are two scenarios in which separate Virtual Chassis merge:

- A Virtual Chassis configuration that had split into two is now merging back into a single configuration because the problem that had caused it to split has been resolved.
- You want to merge two Virtual Chassis that had not previously been configured together.

Every Virtual Chassis configuration has a unique ID that is automatically assigned when the Virtual Chassis configuration is formed. You can also explicitly assign a Virtual Chassis ID using the **set virtual-chassis id** command. A Virtual Chassis ID that you assign takes precedence over automatically assigned Virtual Chassis IDs.

When you reconnect the separate Virtual Chassis configurations or connect them for the first time, the members determine whether or not the separate Virtual Chassis

configurations can merge. The members use the following rules to determine whether a merge is possible:

- If the Virtual Chassis configurations have the same Virtual Chassis ID, then the configurations can merge. If the two Virtual Chassis were formed as the result of a split, they will have the same Virtual Chassis ID.
- If the Virtual Chassis IDs are different, then the two configurations can merge only if both are active (inactive configurations cannot merge, ensuring that members removed from one Virtual Chassis configuration do not become members of another Virtual Chassis configuration). If the configurations to merge are both active and one of them has a user-configured Virtual Chassis ID, this ID becomes the ID of the merged Virtual Chassis. If neither Virtual Chassis has a user-configured Virtual Chassis ID, then the Virtual Chassis ID of the configuration with the highest mastership priority becomes the ID of the merged Virtual Chassis. The resulting merged Virtual Chassis configuration will be active.

When you connect two Virtual Chassis configurations, the following events occur:

1. Connecting the two split Virtual Chassis configurations triggers the shortest-path-first (SPF) algorithm. The SPF algorithm computes the network topology and then triggers the master election algorithm. The master election algorithm waits for the members to synchronize the topology information before running.
2. The master election algorithm merges the Virtual Chassis IDs of all the members.
3. Each member runs the master election algorithm to select a master and a backup from among all members with the same Virtual Chassis IDs. For more information, see “Understanding How the Master in a Virtual Chassis Configuration Is Elected” on page 698.
4. The master determines whether the Virtual Chassis configuration is active or inactive. (See “What Happens When a Virtual Chassis Configuration Splits” on page 712.)
5. If the Virtual Chassis configuration is active, the master assigns roles to all members. If the Virtual Chassis configuration is inactive, the master assigns all members the role of linecard.
6. When the other members receive their role from the master, they change their role to backup or linecard. They also use the active or inactive state information sent by the master to set their own state to active or inactive and to construct the Virtual Chassis member list from the information sent by the master.
7. If the Virtual Chassis state is active, the master waits for messages from the members indicating that they have changed their roles to the assigned roles, and then the master changes its own role to master.

NOTE: When you merge two Virtual Chassis that had not previously been part of the same Virtual Chassis configuration, any configuration settings (such as the settings for Telnet/FTP services, GRES, fast failover, VLANs, and so on) that exist on the new master will become the configuration settings for all members of the new Virtual Chassis, overwriting any other configuration settings.

Related Documentation

- Understanding Virtual Chassis Configuration on page 704
- Example: Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge on page 767
- Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge (CLI Procedure) on page 800
- Disabling Split and Merge in a Virtual Chassis Configuration (CLI Procedure) on page 799

Understanding Automatic Software Update on Virtual Chassis Member Switches

The automatic software update feature automatically updates the Junos OS version on prospective member switches as they are added to a Virtual Chassis configuration of J-EX4200 Ethernet Switches so the new member switch immediately joins the Virtual Chassis configuration and is put in the active state.

For a standalone J-EX4200 switch to join an existing Virtual Chassis configuration, it must be running the same version of Junos OS that is running on the Virtual Chassis master. When the master in a Virtual Chassis configuration detects that a new switch has been added to the configuration, it checks the software version on the new switch. If the software version on the new switch is not the same as the version running on the master, the master keeps the new switch in the inactive state. If you have not enabled the automatic software update feature, you will have to manually install the correct software version on each prospective member switch as it is added to the Virtual Chassis configuration.

Related Documentation

- Understanding Software Upgrade in a Virtual Chassis Configuration on page 698
- Example: Configuring Automatic Software Update on Virtual Chassis Member Switches on page 777
- Configuring Automatic Software Update on Virtual Chassis Member Switches (CLI Procedure) on page 800

CHAPTER 44

Virtual Chassis—Configuration Examples

Example: Configuring a Virtual Chassis with a Master and Backup in a Single Wiring Closet

A Virtual Chassis configuration is a scalable switch. You can provide secure, redundant network accessibility with a basic two-member Virtual Chassis configuration and later expand the Virtual Chassis configuration to provide additional access ports as your office grows.

This example describes how to configure a Virtual Chassis with a master and backup in a single wiring closet:

Requirements

This example uses the following hardware and software components:

- One J-EX4200-48T switch
- One J-EX4200-24T switch
- One uplink module

Before you begin, be sure you have:

1. Rack-mounted the switches. See Mounting a J-EX4200 Switch on Two Posts in a Rack or Cabinet or Mounting a J-EX4200 Switch on Four Posts in a Rack or Cabinet or Mounting a J-EX4200 Switch on a Desk or Other Level Surface.
2. Installed the uplink module. See Installing an Uplink Module in a J-EX4200 Switch.
3. Cabled the switches. See Connecting a Virtual Chassis Cable to a J-EX4200 Switch.

Overview and Topology

A Virtual Chassis configuration allows you to accommodate the networking needs of a growing office. The default configuration of a two-member Virtual Chassis includes a master and a backup switch. In addition to providing more access ports than a single J-EX4200 switch can provide, a Virtual Chassis configuration provides high availability through redundancy.

This example shows a Virtual Chassis configuration composed of two J-EX4200 switches. One of the switches has an uplink module with ports that can be configured to connect to a distribution switch or customer edge (CE) router or that can be configured as Virtual Chassis ports (VCPs) to interconnect with a member switch that is located too far for the dedicated VCP cabling. (The network interfaces on J-EX4200-24F switches can also be configured as VCPs.) For information on configuring the uplink ports as trunk ports to a distribution switch, see “Configuring Gigabit Ethernet Interfaces (CLI Procedure)” on page 919. For an example of configuring uplink ports as VCPs, see “Example: Configuring a Virtual Chassis Interconnected Across Multiple Wiring Closets” on page 733.

By default, after you interconnect the switches with the dedicated VCPs and power on the switches, the VCPs are operational. The mastership priorities and member IDs are assigned by the software. The software elects a master based on several criteria, including how long a member switch has belonged to the Virtual Chassis configuration. For additional details, see “Understanding How the Master in a Virtual Chassis Configuration

Is Elected" on page 698. Therefore, we recommend that you start by powering on only one member switch, the one that you want to function as the master.

NOTE: We recommend that you use the **commit synchronize** command to save any configuration changes that you make to a multimember Virtual Chassis.

The Virtual Chassis configuration provides networking access for 50 onsite workers, who are sitting within range of a single wiring closet. The workers all use personal computers and VoIP phones. As the office grows, you can add more J-EX4200 switches to meet increased needs for access ports.

The topology for this example consists of two switches, one of which contains an uplink module:

- One J-EX4200-48T switch (SWA-0) with 48 access ports, all of which support PoE
- One J-EX4200-24T switch (SWA-1) with 24 access ports, including eight ports that support PoE
- One uplink module, with two 10-Gigabit Ethernet ports, is installed in SWA-1.

Table 102 on page 719 shows the default configuration settings for the two-member Virtual Chassis.

Table 102: Components of the Basic Virtual Chassis Access Switch Topology

Member Switch	Hardware	Member ID	Role and Priority
SWA-0	J-EX4200-48T switch	0	Master: mastership priority 128
SWA-1	J-EX4200-24T switch	1	Backup: mastership priority 128

Figure 14 on page 719 shows that SWA-0 and SWA-1 are interconnected with their dedicated VCPs on the rear panel. The LCD on the front displays the Member ID and Role. SWA-0 also includes an uplink module. Its uplink ports can be used to connect to a distribution switch.

Figure 14: Basic Virtual Chassis with Master and Backup

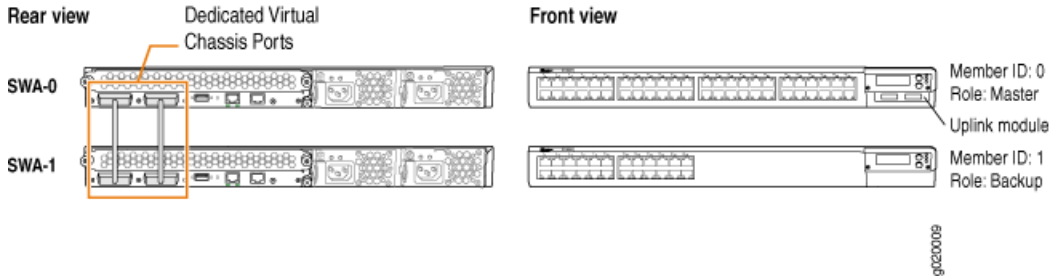

Configuration

Configure a Virtual Chassis with a default master and backup in a single wiring closet:

Step-by-Step Procedure

To configure a Virtual Chassis with master and backup:

1. Make sure the VCPs on the rear panel of the member switches are properly cabled. See Virtual Chassis Cabling Configuration Examples for J-EX4200 Switches.
2. Power on SWA-0 (the member switch that you want to function as the master).
3. Check the front-panel LCD to confirm that the switch has powered on correctly.
4. Run the EZ Setup program on SWA-0, specifying the identification parameters. See “Connecting and Configuring a J-EX Series Switch (CLI Procedure)” on page 161 or “Connecting and Configuring a J-EX Series Switch (J-Web Procedure)” on page 163 for details.
5. Configure SWA-0 with the virtual management Ethernet (VME) interface for out-of-band management of the Virtual Chassis configuration, if desired.

   ```
   [edit]
   user@SWA-0# set interfaces vme unit 0 family inet address /ip-address/mask/
   ```

6. Power on SWA-1.

Verification

To confirm that the Virtual Chassis configuration is operational, perform these tasks:

- Verifying That the Mastership Priority Is Assigned Appropriately on page 720
- Verifying That the VCPs Are Operational on page 721

Verifying That the Mastership Priority Is Assigned Appropriately

Purpose Verify that the master, which has been selected by default, is the member switch that you want to function in that role.

Action

1. Check the front-panel LCD to confirm that the switch has powered on correctly and that a member ID has been assigned.
2. List the member switches of the Virtual Chassis configuration.

```
user@SWA-0> show virtual-chassis status

Virtual Chassis ID: 0019.e250.47a0
                                          Mastership            Neighbor List
Member ID  Status   Serial No    Model    priority    Role      ID  Interface
0 (FPC 0)  Prsnt    AK0207360276 ex4200-48t      128  Master*    1  vcp-0
                                                                 1  vcp-1
1 (FPC 1)  Prsnt    AK0207360281 ex4200-24t      128  Backup     0  vcp-0
                                                                 0  vcp-1

Member ID for next new member: 2 (FPC 2)
```

Meaning The **show virtual-chassis status** command lists the member switches interconnected in a Virtual Chassis configuration with the member IDs that have been assigned by the master, the mastership priority values, and the roles. It also displays the neighbor members with which each member is interconnected. The output shows that SWA-0, member 0, has been assigned default mastership priority 128. Because SWA-0 is the first member to be powered on, it has the most seniority and is therefore assigned the role of master. SWA-1 is powered on after member 0, so it is assigned the role of backup. The member IDs are displayed on the front panel of the switches. Check and confirm whether the default assignment is satisfactory.

Verifying That the VCPs Are Operational

Purpose Verify that the dedicated Virtual Chassis ports interconnecting the switches are operational.

Action Display the Virtual Chassis ports of all the members:

```
user@SWA-0> show virtual-chassis vc-port all-members
fpc0:
--------------------------------------------------------------------------
Interface        Type            Status    Speed        Neighbor
or                                         (mbps)     ID   Interface
PIC / Port
vcp-0            Dedicated        Up        32000      1    vcp-1
vcp-1            Dedicated        Up        32000      1    vcp-0

fpc1:
--------------------------------------------------------------------------
Interface        Type            Status    Speed        Neighbor
or                                         (mbps)     ID   Interface
PIC / Port
vcp-0            Dedicated        Up        32000      1    vcp-0
vcp-1            Dedicated        Up        32000      1    vcp-1
```

Meaning The **show virtual-chassis vc-port** command lists the interfaces that are enabled for the member switches of the Virtual Chassis configuration and shows the status of the interfaces. The output in this example shows that two of the VCPs are operational and two VCPs are not. A single cable has been used to interconnect vcp-0 of member ID 0 and vcp-0 of member ID 1. That interconnection is sufficient for the switch to be operational. However, we recommend that you connect the second set of VCPs for redundancy.

Troubleshooting the Virtual Chassis

To troubleshoot the configuration of a Virtual Chassis, perform these tasks:

Troubleshooting the Assignment of Roles

Problem The master and backup roles are not assigned to the member switches that you want to function in these roles.

Solution Modify the mastership priority values.

To quickly modify the mastership priority of SWA-1 (member ID 1), copy the following command and paste it into the switch terminal window:

```
[edit virtual-chassis]
user@SWA-1# set member 1 mastership-priority 255
```

Troubleshooting the VCPs

Problem The VCPs are down.

Solution

1. Check to make sure that you have cabled the appropriate ports.
2. Check to make sure that the cables are seated properly.

You should generally cable and interconnect both of the VCPs on the member switches, for redundancy and high availability.

Related Documentation

- Example: Expanding a Virtual Chassis Configuration in a Single Wiring Closet on page 722
- Example: Setting Up a Multimember Virtual Chassis Access Switch with a Default Configuration on page 727
- Example: Configuring a Virtual Chassis Using a Preprovisioned Configuration File on page 752
- Configuring a Virtual Chassis (CLI Procedure) on page 781
- Configuring a Virtual Chassis (J-Web Procedure) on page 784

Example: Expanding a Virtual Chassis Configuration in a Single Wiring Closet

A Virtual Chassis configuration is a scalable switch composed of multiple interconnected J-EX4200 switches. Up to ten J-EX4200 switches can be interconnected as a Virtual Chassis configuration.

This example describes how to configure an expanding Virtual Chassis within a single wiring closet:

Requirements

This example uses the following hardware and software components:

- One J-EX4200-48T switch
- Two J-EX4200-24T switches
- One uplink module

Before you begin, be sure you have:

- Confirmed that the existing Virtual Chassis configuration is operating correctly. See “Example: Configuring a Virtual Chassis with a Master and Backup in a Single Wiring Closet” on page 717.

Overview and Topology

A Virtual Chassis configuration can be expanded without disrupting the site's network connectivity. This example describes adding a member switch to an existing Virtual Chassis configuration to provide additional access ports for connecting more PCs and VoIP phones at this location. You can continue to expand the Virtual Chassis configuration with additional members in the same wiring closet, using the same procedure. If you want to expand the Virtual Chassis configuration to include member switches in another wiring closet, see “Example: Configuring a Virtual Chassis Interconnected Across Multiple Wiring Closets” on page 733.

If you want to retain the roles of the existing master and backup switches, explicitly configure the mastership priority of these switches, specifying the highest possible value (255) for both the master and the backup.

During expansion, the existing Virtual Chassis configuration can remain powered on and connected to the network. Before powering up the new switch, interconnect it to the other the switches using the dedicated VCPs on the rear panel. Do not run the EZ Setup program on the added member switch.

This example shows an existing Virtual Chassis configuration composed of two J-EX4200 switches. The Virtual Chassis configuration is being expanded to include a J-EX4200-24T switch as a linecard member.

The topology for this example consists of:

- One J-EX4200-48T switch (SWA-0) with 48 access ports, 8 of which support Power over Ethernet (PoE)
- Two J-EX4200-24T switch (SWA-1 and SWA-2) each with 24 access ports, including 8 ports that support PoE
- One uplink module with two 10-gigabit ports is installed in the J-EX4200-48T switch. These ports can be configured as trunk ports to connect to a distribution switch or customer edge (CE) router or as Virtual Chassis ports (VCPs) to interconnect with a member switch that is located too far for dedicated VCP cabling. (The uplink module ports on the SFP and SFP+ uplink modules and the SFP network interfaces on the J-EX4200-24F switches can also be used for these purposes.) For information on configuring the uplink ports as trunk ports to a distribution switch, see “Configuring Gigabit Ethernet Interfaces (CLI Procedure)” on page 919 or “Configuring Gigabit Ethernet Interfaces (J-Web Procedure)” on page 909. For information on configuring uplink ports as Virtual Chassis ports, see “Setting an Uplink Module Port as a Virtual Chassis Port (CLI Procedure)” on page 792.

Table 103 on page 724 shows the configuration settings for the expanded Virtual Chassis.

Table 103: Components of the Expanded Virtual Chassis Access Switch

Member Switch	Hardware	Member ID	Role in Virtual Chassis
SWA-0	J-EX4200-48T switch	0	master; mastership priority 255
SWA-1	J-EX4200-24T switch	1	backup; mastership priority 255
SWA-2	J-EX4200-24T switch	2	linecard; mastership priority 128

Figure 15 on page 724 shows that the three member switches (SWA-0, SWA-1 and SWA-2) are interconnected with their dedicated VCPs on the rear panel. The LCD on the front displays the member ID and role. SWA-0 also includes an uplink module. Its uplink ports can be used to connect to a distribution switch.

Figure 15: Expanded Virtual Chassis in Single Wiring Closet

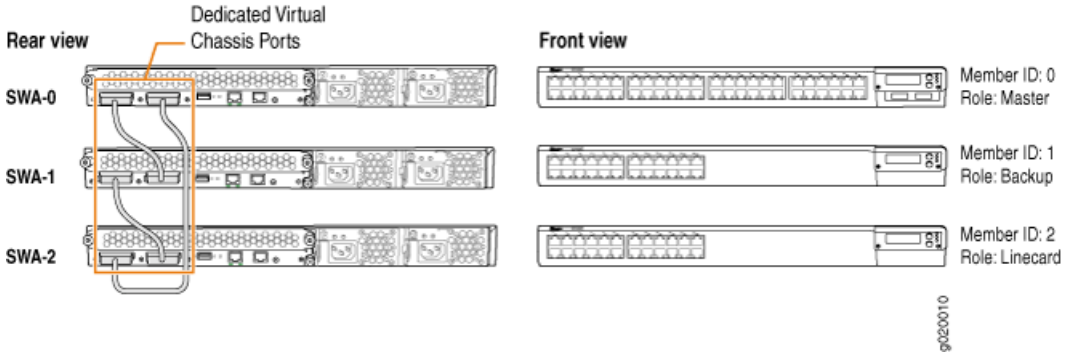

Configuration

To expand a Virtual Chassis configuration to include additional member switches within a single wiring closet, perform these tasks:

NOTE: We recommend that you use the **commit synchronize** command to save any configuration changes that you make to a multimember Virtual Chassis configuration.

CLI Quick Configuration

To maintain the master and backup roles of the existing members and ensure that the new member switch functions in a linecard role, copy the following commands and paste them into the terminal window:

```
[edit]
user@SWA-0# set virtual-chassis member 0 mastership-priority 255
user@SWA-1# set virtual-chassis member 1 mastership-priority 255
```

Step-by-Step Procedure

To ensure that the existing member switches retain their current roles and to add another member switch in a linecard role:

1. Configure the mastership priority of SWA-0 (member 0) to be the highest possible value, thereby ensuring that it functions as the master of the expanded Virtual Chassis configuration.

   ```
   [edit virtual-chassis]
   user@SWA-0# set member 0 mastership-priority 255
   ```

2. Configure the mastership priority of SWA-1 (member 1) to be the highest possible value. This setting is recommended for high availability and smooth transition of mastership in case the original master becomes unavailable.

   ```
   [edit virtual-chassis]
   user@SWA-1# set member 1 mastership-priority 255
   ```

3. Interconnect the unpowered SWA-2 with SWA-0 and SWA-1 using the dedicated VCPs on the rear panel. See Virtual Chassis Cabling Configuration Examples for J-EX4200 Switches for additional information.

4. Power on SWA-2.

   You do not need to configure or run EZ Setup on SWA-2. The identification parameters that were set up for the master apply implicitly to all members of the Virtual Chassis configuration. SWA-2 functions in a linecard role, since SWA-0 and SWA-1 have been configured to the highest mastership priority values.

Verification

To verify that the new switch has been added as a linecard and that its VCPs are operational, perform these tasks:

- Verifying That the New Switch Has Been Added as a Linecard on page 725
- Verifying That the VCPs Are Operational on page 726

Verifying That the New Switch Has Been Added as a Linecard

Purpose Verify that SWA-2 has been added in a linecard role to the Virtual Chassis configuration.

Action Use the **show virtual-chassis status** command to list the member switches with their member IDs, mastership priority values, and assigned roles.

```
user@SWA-0> show virtual-chassis status

Virtual Chassis ID: 0000.e255.00e0

                                          Mastership            Neighbor List
Member ID  Status  Serial No   Model     Priority  Role      ID Interface

0 (FPC 0)  Prsnt   abc123      ex4200-48t    255  Master*    1 vcp-0
                                                             2 vcp-1

1 (FPC 1)  Prsnt   def456      ex4200-24t    255  Backup     2 vcp-0
                                                             0 vcp-1
```

```
2 (FPC 2)  Prsnt   abd231      ex4200-24tp    128  Linecard    0 vcp-0
                                                               1 vcp-1
```

Meaning The **show virtual-chassis status** command lists the member switches of the Virtual Chassis configuration with the member IDs and mastership priority values. It also displays the neighbor members with which each member is interconnected. This output shows that SWA-2 has been assigned member ID 2 and has the default mastership priority value 128. Because the mastership priority is lower than the mastership priority of the other members, SWA-2 functions in the linecard role. You can continue to add more member switches, following the same procedure. It is possible to have multiple members in linecard roles with the same mastership priority value.

Verifying That the VCPs Are Operational

Purpose Verify that the dedicated VCPs interconnecting the member switches are operational.

Action List the VCP interfaces on the Virtual Chassis configuration.

```
user@SWA-0>show virtual-chassis vc-port all-members
fpc0:
--------------------------------------------------------------------------
Interface        Type             Status
or
PIC / Port
vcp-0            Dedicated        Up
vcp-1            Dedicated        Up

fpc1:
--------------------------------------------------------------------------
Interface        Type             Status
or
PIC / Port
vcp-0            Dedicated        Up
vcp-1            Dedicated        Up

fpc2:
--------------------------------------------------------------------------
Interface        Type             Status
or
PIC / Port
vcp-0            Dedicated        Up
vcp-1            Dedicated        Up
```

Meaning The **show virtual-chassis vc-port all-members** command lists all the interfaces for the Virtual Chassis configuration. In this case, no VCP uplinks have been configured. However, the VCP interfaces are automatically configured and enabled when you interconnect member switches using the dedicated Virtual Chassis ports. There are two dedicated VCPs on the rear panel of each J-EX4200 switch. It is recommended that you interconnect the member switches using both VCPs for redundancy. The VCP interfaces are identified simply as **vcp-0** and **vcp-1**. The **fpc** number is the same as the member ID.

Troubleshooting

To troubleshoot the configuration of an expanded Virtual Chassis, perform these tasks:

Troubleshooting Mastership Priority

Problem You want to designate a different member as the master.

Solution Change the mastership priority value or values of the switches, designating the highest mastership priority value for the switch that you want to be master.

1. Lower the mastership priority of the existing master (member 0).

```
[edit virtual-chassis]
user@SWA-0# set member 0 mastership-priority 1
```

2. Set the mastership priority of the member that you want to be the master to the highest possible value (255):

```
[edit virtual-chassis]
user@SWA-2# set member 2 mastership-priority 255
```

Troubleshooting Nonoperational VCPs

Problem The VCP interface shows a status of **down**.

Solution Check the cable to make sure that it is properly and securely connected to the VCPs.

Related Documentation

- Example: Setting Up a Multimember Virtual Chassis Access Switch with a Default Configuration on page 727
- Configuring a Virtual Chassis (CLI Procedure) on page 781
- Configuring a Virtual Chassis (J-Web Procedure) on page 784

Example: Setting Up a Multimember Virtual Chassis Access Switch with a Default Configuration

You can configure a multimember Virtual Chassis access switch in a single wiring closet without setting any parameters—by simply cabling the switches together, using the dedicated Virtual Chassis ports (VCPs). You do not need to modify the default configuration to enable these ports. They are operational by default. The Virtual Chassis configuration automatically assigns the master, backup, and linecard roles, based on the sequence in which the switches are powered on and other factors in the master election algorithm. See “Understanding How the Master in a Virtual Chassis Configuration Is Elected” on page 698.

TIP: We recommend that you explicitly configure the mastership priority of the switches to ensure that the switches continue to perform the desired roles when additional switches are added or other changes occur. However, it is possible to use the default configuration described in this example.

This example describes how to configure a multimember Virtual Chassis in a single wiring closet, using the default role assignments:

Requirements

This example uses the following hardware and software components:

- Two J-EX4200-48T switches
- Four J-EX4200-24T switches

Overview and Topology

A Virtual Chassis configuration is easily expandable. This example shows a Virtual Chassis configuration composed of six J-EX4200 switches. It provides networking access for 180 onsite workers, who are sitting within range of a single wiring closet. The six combined switches are identified by a single host name and managed through a global management IP address.

To set up a multimember Virtual Chassis configuration within a single wiring closet, you need to run the EZ Setup program only once. Connect to the master and run EZ Setup to specify its identification, time zone, and network properties. When additional switches are connected through the Virtual Chassis ports (VCPs), they automatically receive the same properties that were specified for the master.

The topology for this example (see Figure 16 on page 729) consists of six switches:

- Two J-EX4200-48T switches (SWA-0 and SWA-1) with 48 access ports, 8 of which support Power over Ethernet (PoE)
- Four J-EX4200-24T switches (SWA-2, SWA-3, SWA-4, and SWA-5) with 24 access ports, 8 of which support PoE

Figure 16 on page 729 shows that all the member switches are interconnected with the dedicated VCPs on the rear panel. The LCD on the front displays the member ID and role.

Figure 16: Default Configuration of Multimember Virtual Chassis in a Single Wiring Closet

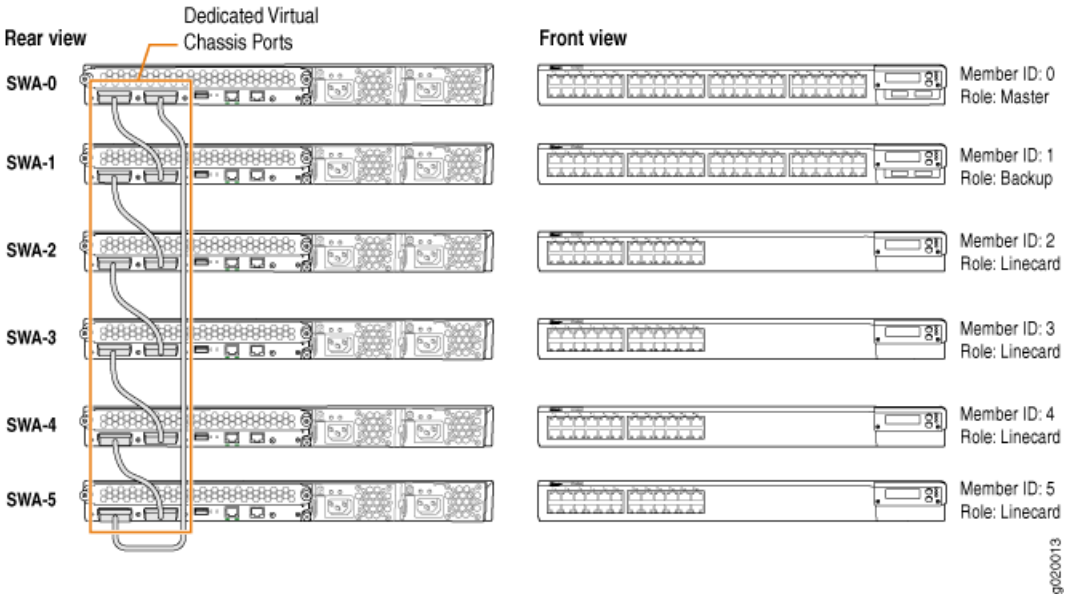

Configuration

Configure a multimember Virtual Chassis access switch in a single wiring closet using the factory defaults:

CLI Quick Configuration

By default, after you interconnect the switches with the dedicated VCPs and power on the switches, the VCPs are operational. The mastership priorities and member IDs are assigned by the software. To determine which switch has been selected as the master, check the LCD on the front panel. It should be the first switch that you power on. The backup should be the second switch that you power on. The other switches are all linecards. Wait at least one minute after powering on the master, before continuing to power on the other switches.

Step-by-Step Procedure

To configure a multimember Virtual Chassis with default role assignments:

1. Make sure the dedicated VCPs on the rear panel are properly cabled. See Virtual Chassis Cabling Configuration Examples for J-EX4200 Switches for additional information.
2. Power on the switch that you want to function as the master (SWA-0). This examples uses one of the larger switches (J-EX4200-48T) as the master.
3. Check the front panel LCD to confirm that the switch has powered on correctly and that a member ID has been assigned.
4. Run the EZ Setup program on SWA-0, the master, specifying the identification parameters. See “Connecting and Configuring a J-EX Series Switch (CLI Procedure)” on page 161 or “Connecting and Configuring a J-EX Series Switch (J-Web Procedure)” on page 163 for details.

5. Configure SWA-0 with the virtual management Ethernet (VME) interface for out-of-band management of the Virtual Chassis configuration, if desired.

```
[edit]
user@SWA-0# set interfaces vme unit 0 family inet address /ip-address/mask/
```

6. After a lapse of at least one minute, power on SWA-1. This example uses the second J-EX4200-48T switch as the backup.
7. Check the front panel LCD to confirm that the switch has powered on correctly and that a member ID has been assigned.
8. Power on SWA-2, and check the front panels to make sure that the switch is operating correctly.
9. Continue to power on the member switches one by one, checking the front panels as you proceed.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying the Member IDs and Roles of the Member Switches on page 730
- Verifying That the VCPs Are Operational on page 731

Verifying the Member IDs and Roles of the Member Switches

Purpose Verify that all the interconnected member switches are included within the Virtual Chassis configuration and that their roles are assigned appropriately.

Action Display the members of the Virtual Chassis configuration:

```
user@SWA-0> show virtual-chassis status

Virtual Chassis ID: 0000.e255.00e0

                                          Mastership            Neighbor List
Member ID  Status  Serial No   Model     Priority  Role        ID Interface
0 (FPC 0)  Prsnt   abc123      ex4200-48t    128  Master*     1 vcp-0
                                                              5 vcp-1
1 (FPC 1)  Prsnt   def123      ex4200-48p    128  Backup      2 vcp-0
                                                              0 vcp-1
2 (FPC 2)  Prsnt   abd231      ex4200-24t    128  Linecard    3 vcp-0
                                                              1 vcp-1
3 (FPC 3)  Prsnt   cab123      ex4200-24t    128  Linecard    4 vcp-0
                                                              2 vcp-1
4 (FPC 4)  Prsnt   fed456      ex4200-24t    128  Linecard    5 vcp-0
                                                              3 vcp-1
5 (FPC 5)  Prsnt   jkl231      ex4200-24t    128  Linecard    0 vcp-0
                                                              4 vcp-1
```

Meaning The **show virtual-chassis status** command lists the member switches of the Virtual Chassis configuration with the member IDs and mastership priority values. It also displays the neighbor members with which each member is interconnected. The **fpc** number is the same as the member ID.

Verifying That the VCPs Are Operational

Purpose Verify that the dedicated VCPs interconnecting the member switches are operational.

Action Display the Virtual Chassis interfaces.

```
user@SWA-0> show virtual-chassis vc-port all-members
 fpc0:
--------------------------------------------------------------------------------
Interface        Type             Status
or
PIC / Port
vcp-0            Dedicated        Up
vcp-1            Dedicated        Up

fpc1:
--------------------------------------------------------------------------------
Interface        Type             Status
or
PIC / Port
vcp-0            Dedicated        Up
vcp-1            Dedicated        Up

fpc2:
--------------------------------------------------------------------------------
Interface        Type             Status
or
PIC / Port
vcp-0            Dedicated        Up
vcp-1            Dedicated        Up

fpc3:
--------------------------------------------------------------------------------
Interface        Type             Status
or
PIC / Port
vcp-0            Dedicated        Up
vcp-1            Dedicated        Up

fpc4:
--------------------------------------------------------------------------------
Interface        Type             Status
or
PIC / Port
vcp-0            Dedicated        Up
vcp-1            Dedicated        Up

fpc5:

Interface        Type             Status
or
PIC / Port
```

```
vcp-0           Dedicated       Up
vcp-1           Dedicated       Up
```

Meaning The **show virtual-chassis vc-port all-members** command lists the Virtual Chassis interfaces that are enabled for the member switches of the Virtual Chassis configuration and shows the status of the interfaces. In this case, no VCP uplinks have been configured. However, the VCP interfaces are automatically configured and enabled when you interconnect member switches using the dedicated VCPs. There are two dedicated VCPs on the rear panel of each J-EX4200 switch. The dedicated VCP interfaces are identified simply as vcp-0 and vcp-1. They do not use the standard interface address (in which the member ID is represented by the first digit). The output in this example shows that all interfaces are operational. The **fpc** number is the same as the member ID.

Troubleshooting

To troubleshoot the configuration of a multimember Virtual Chassis in a single wiring closet, perform these tasks:

Troubleshooting Mastership Priority

Problem You want to explicitly designate one member as the master and another as backup.

Solution Change the mastership priority value of the member that you want to function as master, designating the highest mastership priority value that member.

NOTE: These configuration changes are made through the current master, SWA-0.

1. Configure mastership priority of member 0 to be the highest possible value.

```
[edit virtual-chassis]
user@SWA-0# set member 0 mastership-priority 255
```

2. Set the mastership priority of another member that you want to function as the backup member as the same value:

```
[edit virtual-chassis]
user@SWA-0# set member 2 mastership-priority 255
```

Troubleshooting Nonoperational VCPs

Problem The VCP interface shows a status of **down**.

Solution Check the cable to make sure that it is properly and securely connected to the VCPs.

Related Documentation

- Example: Configuring a Virtual Chassis with a Master and Backup in a Single Wiring Closet on page 717
- Example: Configuring a Virtual Chassis Interconnected Across Multiple Wiring Closets on page 733

Example: Configuring a Virtual Chassis Interconnected Across Multiple Wiring Closets

A Virtual Chassis configuration is a very adaptable access switch solution. You can install member switches in different wiring closets, interconnecting the member switches by cabling and configuring uplink module ports and SFP network ports on J-EX4200-24F switches as Virtual Chassis ports (VCPs).

This example shows how to use uplink VCPs to connect Virtual Chassis members that are located too far apart to be connected using the dedicated VCPs. Uplink VCPs can also be used to connect Virtual Chassis members to form link aggregation groups (LAGs). For the latter usage, see “Example: Configuring Link Aggregation Groups Using Uplink Virtual Chassis Ports” on page 769.

NOTE: You can also configure the SFP networks ports on J-EX4200-24F switches as VCPs to connect Virtual Chassis member switches across wiring closets and to form LAGs.

This example describes how to configure a Virtual Chassis access switch interconnected across wiring closets:

Requirements

This example uses the following hardware and software components:

- Four J-EX4200 switches
- Four uplink modules

Before you interconnect the members of the Virtual Chassis configuration across wiring closets, be sure you have:

1. Installed an uplink module in each member switch. See Installing an Uplink Module in a J-EX4200 Switch.
2. Powered on, connected, and run the EZSetup program on SWA-0 (see Table 104 on page 735 for switch names used in this example). See "Connecting and Configuring a J-EX Series Switch (CLI Procedure)" on page 161 or "Connecting and Configuring a J-EX Series Switch (J-Web Procedure)" on page 163 for details.
3. Configured SWA-0 with the virtual management Ethernet (VME) interface for remote, out-of-band management of the Virtual Chassis configuration, if desired. See "Configuring the Virtual Management Ethernet Interface for Global Management of a Virtual Chassis (CLI Procedure)" on page 797.
4. Interconnected SWA-0 and SWA-1 using the dedicated VCPs on the rear panel. SWA-1 must not be powered on at this time.
5. Interconnected SWA-2 and SWA-3 using the dedicated VCPs on the rear panel. SWA-2 and SWA-3 must not be powered on at this time.

Overview and Topology

In this example, four J-EX4200 switches will be interconnected in a Virtual Chassis configuration. Two of these (SWA-0 and SWA-1) are located in wiring closet A and the two other (SWA-2 and SWA-3) are located in wiring closet B.

For ease of monitoring and manageability, we want to interconnect all four switches as members of a Virtual Chassis configuration. Prior to configuring the Virtual Chassis, we installed uplink modules in each of the member switches. In this example, uplink modules are installed in all four members so that there are redundant VCP connections across the wiring closets. If you want to expand this configuration to include more members within these wiring closets, you do not need to add any more uplink modules. Simply use the dedicated VCPs on the rear panel. The redundancy of uplink VCPs provided in this example is sufficient.

We have interconnected the switches in wiring closet A and also interconnected the ones in wiring closet B using the dedicated VCPs. The interfaces for the dedicated VCPs are operational by default. They do not need to be configured.

However, the Virtual Chassis cables that interconnect the dedicated VCPs of member switches within a single wiring closet are not long enough to connect member switches across wiring closets. Instead, we will use the fiber-optic cable connections in the uplink modules to interconnect the member switches in wiring closet A to the member switches in wiring closet B. You only need to interconnect one member switch in wiring closet A to one in wiring closet B to form the Virtual Chassis configuration. However, for redundancy, this example connects uplink module ports from the two member switches in wiring closet A to the two member switches in wiring closet B.

We will specify the highest mastership priority value (255) for SWA-0 to make it the master before we power on SWA-1. Because SWA-0 and SWA-1 are interconnected with

the dedicated VCPs, the master detects that SWA-1 is a member of its Virtual Chassis configuration and assigns it a member ID.

We configure SWA-2 in wiring closet B without running EZSetup by directly connecting to the console port. If you wish, you can run EZSetup and specify identification parameters. Later, when you interconnect SWA-2 with SWA-0, the master of the Virtual Chassis configuration, the master overwrites any conflicting parameters.

We will use SWA-2 as the backup of the Virtual Chassis configuration. If a problem occurs in wiring closet A, SWA-2 would take control of the Virtual Chassis configuration and maintain the network connections. We will configure the same mastership priority value for SWA-2 (255) that we configured for the master. Because we power on SWA-0 before we power on SWA-2, SWA-0 has additional prioritization properties that allow it to retain mastership of the Virtual Chassis configuration. See “Understanding How the Master in a Virtual Chassis Configuration Is Elected” on page 698. We recommend setting identical mastership priority values for the master and backup members for high availability and smooth transition of mastership in case the original master becomes unavailable. (Setting identical mastership priority values for the master and backup members prevents the previous master from pre-empting the master role from the new master when the previous master comes back online.)

After we have configured SWA-2 and set one of its uplink module ports as an uplink VCP, we will interconnect its uplink VCP with an uplink VCP on SWA-0.

Finally, we will power on SWA-3. Because SWA-3 is interconnected with SWA-2 using the dedicated VCPs on the rear panel, the master will detect that SWA-3 is part of the expanded Virtual Chassis configuration and assign it member ID **3**. For redundancy, we will configure an uplink VCP on SWA-3 through the master and interconnect that uplink VCP with an uplink VCP on SWA-1.

Table 104 on page 735 shows the Virtual Chassis configuration settings for a Virtual Chassis composed of member switches in different wiring closets.

Table 104: Components of a Virtual Chassis Interconnected Across Multiple Wiring Closets

Switch	Member ID	Role and Priority	Location
SWA-0	0	master; mastership priority 255	Wiring closet A
SWA-1	1	linecard; mastership priority 128	Wiring closet A
SWA-2	2	backup; mastership priority 255	Wiring closet B
SWA-3	3	linecard; mastership priority 128	Wiring closet B

Figure 17 on page 736 shows the different types of interconnections used for this Virtual Chassis configuration. The rear view shows the member switches within each wiring

closet interconnected to each other using the dedicated VCPs. The front view shows the uplink VCPs interconnected across the wiring closets.

Figure 17: Virtual Chassis Interconnected Across Wiring Closets

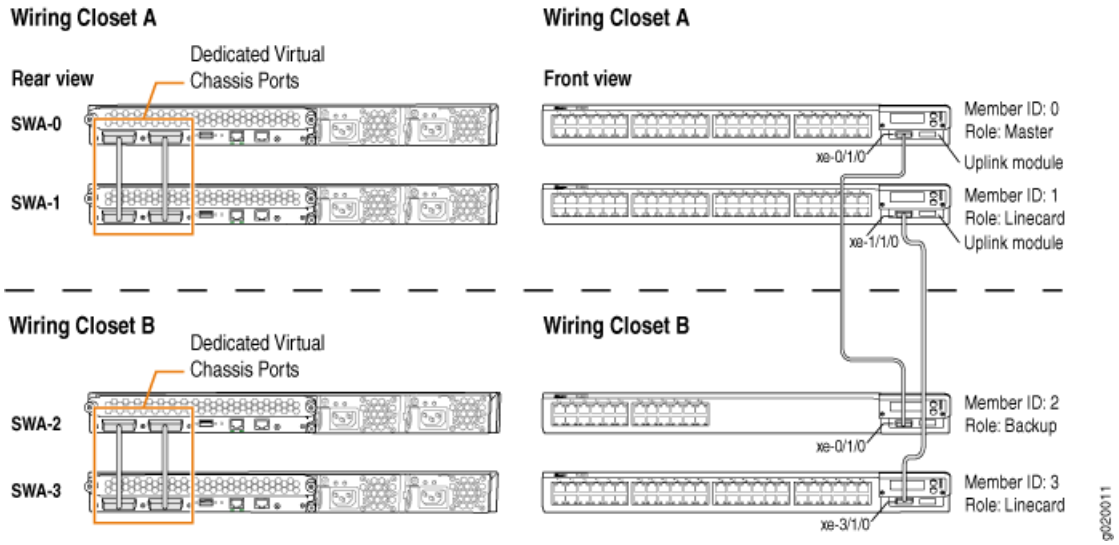

Configuration

To configure the Virtual Chassis across multiple wiring closets, perform these tasks:

Step-by-Step Procedure

To configure a Virtual Chassis across multiple wiring closets:

1. Configure the mastership priority of SWA-0 (member **0**) to be the highest possible value (**255**), thereby ensuring that it functions as the master of the expanded Virtual Chassis configuration:

   ```
   [edit virtual-chassis]
   user@SWA-0# set member 0 mastership-priority 255
   ```

2. Prepare the members in wiring closet A for interconnecting with the member switches in wiring closet B by setting uplink VCPs for member 0 and member 1:

   ```
   user@SWA-0> request virtual-chassis vc-port set pic-slot 1 port 0
   user@SWA-0> request virtual-chassis vc-port set pic-slot 1 port 0 member 1
   ```

   NOTE:
   - For redundancy, this example configures an uplink VCP in both SWA-0 and SWA-1.
   - This example omits the specification of the member *member-id* option in configuring an uplink VCP for SWA-0 (and, later, for SWA-2). The command applies by default to the switch where it is executed.

3. Prepare SWA-2 in wiring closet B for interconnecting with the Virtual Chassis configuration by configuring its mastership priority to be the highest possible value (255). Its member ID is currently **0**, because it is not yet interconnected with the

other members of the Virtual Chassis configuration. It is operating as a standalone switch. Its member ID will change when it is interconnected.

```
[edit virtual-chassis]
user@SWA-2# set member 0 mastership-priority 255
```

NOTE: SWA-2 is configured with the same mastership priority value that we configured for SWA-0. However, the longer uptime of SWA-0 ensures that, once the interconnection is made, SWA-0 functions as the master and SWA-2 functions as the backup.

4. Specify one uplink module port in SWA-2 as an uplink VCP. Its member ID is **0**, because it is not yet interconnected with the other members of the Virtual Chassis configuration.

NOTE: The setting of the uplink VCP remains intact when SWA-2 reboots and joins the Virtual Chassis configuration as member 2.

```
user@SWA-2> request virtual-chassis vc-port set pic-slot 1 port 0
```

5. Physically interconnect SWA-0 and SWA-2 across wiring closets using their uplink VCPs. Although SWA-0 and SWA-2 have the same mastership priority value (255), SWA-0 was powered on first and thus has longer uptime. This results in SWA-0 retaining mastership while SWA-2 reboots and joins the now expanded Virtual Chassis configuration as the backup, with member ID **2**.

6. Power on SWA-3. It joins the expanded Virtual Chassis configuration as member **3**.

NOTE: Member ID 3 is assigned to SWA-3 is 3, because SWA-3 was powered on after members 0, 1, and 2.

7. Because SWA-3 is now interconnected as a member of the Virtual Chassis configuration, you can specify a redundant uplink VCP on SWA-3 through the master of the Virtual Chassis configuration:

```
user@SWA-0> request virtual-chassis vc-port set pic-slot 1 port 0 member 3
```

8. Physically interconnect SWA-3 and SWA-1 across wiring closets using their uplink VCPs. Both SWA-1 and SWA-3 have the default mastership priority value (128) and function in a linecard role.

NOTE: We recommend that you use the **commit synchronize** command to save any configuration changes that you make to a multimember Virtual Chassis.

Results Display the results of the configuration on SWA-0:

```
[edit]
user@SWA-0# show virtual-chassis
  member 0 {
    mastership-priority 255;
  }
  member 1 {
    mastership-priority 128;
  }
  member 2 {
    mastership-priority 255;
  }
  member 3 {
    mastership-priority 128;
  }
}
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying the Member IDs and Roles of the Member Switches on page 738
- Verifying that the Dedicated VCPs and Uplink VCPs Are Operational on page 739

Verifying the Member IDs and Roles of the Member Switches

Purpose Verify that all the interconnected member switches are included within the Virtual Chassis configuration and that their roles are assigned appropriately.

Action Display the members of the Virtual Chassis configuration:

```
user@SWA-0> show virtual-chassis status

Virtual Chassis ID: 0000.e255.00e0

                                          Mastership            Neighbor List
Member ID  Status  Serial No   Model     Priority  Role        ID Interface

0 (FPC 0)  Prsnt   abc123      ex4200-48t    255   Master*     1 vcp-0
                                                               2 vcp-1
                                                               2 vcp-255/1/0

1 (FPC 1)  Prsnt   def456      ex4200-24t    128   Linecard    0 vcp-0
                                                               0 vcp-1
                                                               3 vcp-255/1/0

2 (FPC 2)  Prsnt   ghi789      ex4200-48t    255   Backup      3 vcp-0

                                                                3 vcp-1
                                                               0 vcp-255/1/0

3 (FPC 3)  Prsnt   jkl012      ex4200-24t    128   Linecard    2 vcp-0
                                                               2 vcp-1
                                                               3 vcp-255/1/0
```

Meaning The **show virtual-chassis status** command lists the member switches interconnected as a Virtual Chassis configuration with the member IDs that have been assigned by the

master, the mastership priority values, and the roles. It also displays the neighbor members with which each member is interconnected.

Verifying that the Dedicated VCPs and Uplink VCPs Are Operational

Purpose Verify that the dedicated VCPs interconnecting member switches in wiring closet A and the uplink VCPs interconnecting the member switches between wiring closets are operational.

Action Display the Virtual Chassis interfaces:

```
user@SWA-0> show virtual-chassis status  all-members

fpc0:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status       Speed        Neighbor
or                             ID                (mbps)       ID  Interface
PIC / Port
vcp-0       Dedicated           1    Up           32000
vcp-1       Dedicated           2    Up           32000        1   vcp-0
1/0         Auto-Configured    -1    Up           1000         2   vcp-255/1/0

fpc1:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status       Speed        Neighbor
or                             ID                (mbps)       ID  Interface
PIC / Port
vcp-0       Dedicated           1    Up           32000        0   vcp-0
vcp-1       Dedicated           2    Up           32000        0   vcp-1
1/0         Auto-Configured    -1    Up           1000         3   vcp-255/1/0

fpc2:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status       Speed        Neighbor
or                             ID                (mbps)       ID  Interface
PIC / Port
vcp-0       Dedicated           1    Up           32000        3   vcp-0
vcp-1       Dedicated           2    Up           32000
1/0         Auto-Configured    -1    Up           1000         0   vcp-255/1/0

fpc3:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status       Speed        Neighbor
or                             ID                (mbps)       ID  Interface
PIC / Port
vcp-0       Dedicated           1    Up           32000        2   vcp-0
vcp-1       Dedicated           2    Up           32000        2   vcp-1
1/0         Auto-Configured    -1    Up           1000         1   vcp-255/1/0
```

Meaning The dedicated VCPs are displayed as **vcp-0** and **vcp-1**. The interface on the switch that has been set as an uplink VCP is displayed as **1/0**. The member interface names of uplink VCPs are of the form **vcp-255/*pic/port***—for example, **vcp-255/1/0**. In that name, **vcp-255** indicates that the interface is an uplink VCP, **1** is the uplink PIC number, and **0** is the uplink port number. The **fpc** number is the same as the member ID. The **Trunk ID** is a positive

number ID assigned to the LAG formed by the Virtual Chassis. If no LAG is formed, the value is **–1**.

Troubleshooting

To troubleshoot a Virtual Chassis configuration that is interconnected across wiring closets, perform these tasks:

Troubleshooting Nonoperational VCPs

Problem A uplink VCP shows a status of **down**.

Solution

- Check the cable to make sure that it is properly and securely connected to the ports.
- If the VCP is an uplink module port, make sure that it has been explicitly set as an uplink VCP.
- If the VCP is an uplink module port, make sure that you have specified the options (***pic-slot***, ***port***, and ***member***) correctly.

Related Documentation

- Example: Configuring a Virtual Chassis with a Master and Backup in a Single Wiring Closet on page 717
- Example: Expanding a Virtual Chassis Configuration in a Single Wiring Closet on page 722
- Example: Setting Up a Multimember Virtual Chassis Access Switch with a Default Configuration on page 727
- Setting an Uplink Module Port as a Virtual Chassis Port (CLI Procedure) on page 792

Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch

J-EX Series switches allow you to combine multiple Ethernet links into one logical interface for higher bandwidth and redundancy. The ports that are combined in this manner are referred to as a link aggregation group (LAG) or bundle. The number of Ethernet links you can combine into a LAG depends on your J-EX Series switch model. See “Understanding Aggregated Ethernet Interfaces and LACP” on page 867 for more information.

This example describes how to configure uplink LAGs to connect a Virtual Chassis access switch to a Virtual Chassis distribution switch:

- Requirements on page 741
- Overview and Topology on page 741
- Configuration on page 743
- Verification on page 745
- Troubleshooting on page 746

Requirements

This example uses the following software and hardware components:

- Two J-EX4200-48T switches
- Two J-EX4200-24F switches
- Four uplink modules

Before you configure the LAGs, be sure you have:

- Configured the Virtual Chassis switches. See “Example: Configuring a Virtual Chassis with a Master and Backup in a Single Wiring Closet” on page 717.
- Configured the uplink ports on the switches as trunk ports. See “Configuring Gigabit Ethernet Interfaces (CLI Procedure)” on page 919.

Overview and Topology

For maximum speed and resiliency, you can combine uplinks between an access switch and a distribution switch into LAGs. Using LAGs can be particularly effective when connecting a multimember Virtual Chassis access switch to a multimember Virtual Chassis distribution switch.

The Virtual Chassis access switch in this example is composed of two member switches. Each member switch has an uplink module with two 10-Gigabit Ethernet ports. These ports are configured as trunk ports, connecting the access switch with the distribution switch.

Configuring the uplinks as LAGs has the following advantages:

- Link Aggregation Control Protocol (LACP) can optionally be configured for link negotiation.
- It doubles the speed of each uplink from 10 Gbps to 20 Gbps.
- If one physical port is lost for any reason (a cable is unplugged or a switch port fails, or one member switch is unavailable), the logical port transparently continues to function over the remaining physical port.

The topology used in this example consists of one Virtual Chassis access switch and one Virtual Chassis distribution switch. The access switch is composed of two J-EX4200-48T switches (SWA-0 and SWA-1), interconnected to each other with their Virtual Chassis ports (VCPs) as member switches of Host-A. The distribution switch is composed of two J-EX4200-24F switches (SWD-0 and SWD-1), interconnected with their VCPs as member switches of Host-D.

Each member of the access switch has an uplink module installed. Each uplink module has two ports. The uplinks are configured to act as trunk ports, connecting the access switch with the distribution switch. One uplink port from SWA-0 and one uplink port from SWA-1 are combined as LAG **ae0** to SWD-0. This link is used for one VLAN. The remaining uplink ports from SWA-0 and from SWA-1 are combined as a second LAG connection (**ae1**) to SWD-1. LAG **ae1** is used for another VLAN.

NOTE: If the remote end of the LAG link is a security device, LACP might not be supported because security devices require a deterministic configuration. In this case, do not configure LACP. All links in the LAG are permanently operational unless the switch detects a link failure within the Ethernet physical layer or data link layers.

Figure 18: Topology for LAGs Connecting a Virtual Chassis Access Switch to a Virtual Chassis Distribution Switch

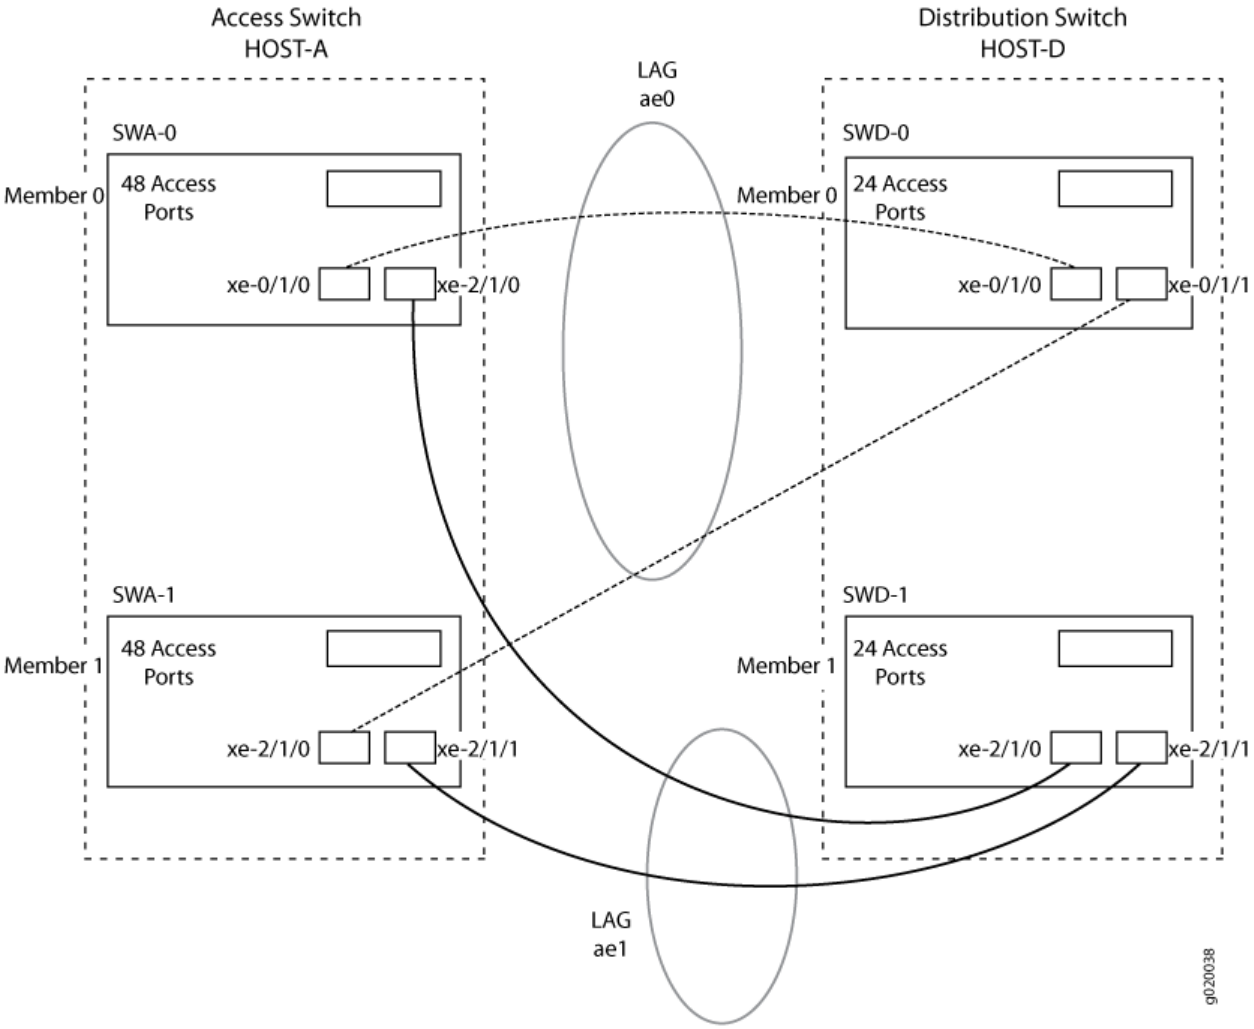

Table 105 on page 743 details the topology used in this configuration example.

Table 105: Components of the Topology for Connecting Virtual Chassis Access Switches to a Virtual Chassis Distribution Switch

Switch	Hostname and VCID	Base Hardware	Uplink Module	Member ID	Trunk Port
SWA-0	Host-A Access switch VCID 1	J-EX4200-48T switch	One uplink module	0	**ge-0/1/0** to SWD-0 **ge-0/1/1** to SWD-1
SWA-1	Host-A Access switch VCID 1	J-EX4200-48T switch	One uplink module	1	**ge-1/1/0** to SWD-0 **ge-1/1/1** to SWD-1
SWD-0	Host-D Distribution switch VCID 4	J-EX4200 L-24F switch	One uplink module	0	**ge-0/1/0** to SWA-0 **ge-0/1/1** to SWA-1
SWD-1	Host-D Distribution switch VCID 4	J-EX4200 L-24F switch	One uplink module	1	**ge-1/1/0** to SWA-0 **xge-1/1/1** to SWA-1

Configuration

To configure two uplink LAGs from the Virtual Chassis access switch to the Virtual Chassis distribution switch:

CLI Quick Configuration

To quickly configure aggregated Ethernet high-speed uplinks between a Virtual Chassis access switch and a Virtual Chassis distribution switch, copy the following commands and paste them into the switch terminal window:

```
[edit]
set chassis aggregated-devices ethernet device-count 2
set interfaces ae0 aggregated-ether-options minimum-links 2
set interfaces ae0 aggregated-ether-options link-speed 10g
set interfaces ae1 aggregated-ether-options minimum-links 2
set interfaces ae1 aggregated-ether-options link-speed 10g
set interfaces ae0 unit 0 family inet address 192.0.2.0/25
set interfaces ae1 unit 1 family inet address 192.0.2.128/25
set interfaces ge-0/1/0 ether-options 802.ad ae0
set interfaces ge-1/1/0 ether-options 802.ad ae0
set interfaces ge-0/1/1 ether-options 802.ad ae1
set interfaces ge-1/1/1 ether-options 802.ad ae1
```

Step-by-Step Procedure

To configure aggregated Ethernet high-speed uplinks between a Virtual Chassis access switch and a Virtual Chassis distribution switch:

1. Specify the number of LAGs to be created on the chassis:

   ```
   [edit chassis]
   user@Host-A# set aggregated-devices ethernet device-count 2
   ```

2. Specify the number of links that need to be present for the **ae0** LAG interface to be **up**:

```
[edit interfaces]
user@Host-A# set ae0 aggregated-ether-options minimum-links 2
```

3. Specify the number of links that need to be present for the **ae1** LAG interface to be **up**:

```
[edit interfaces]
user@Host-A# set ae1 aggregated-ether-options minimum-links 2
```

4. Specify the media speed of the **ae0** link:

```
[edit interfaces]
user@Host-A# set ae0 aggregated-ether-options link-speed 10g
```

5. Specify the media speed of the **ae1** link:

```
[edit interfaces]
user@Host-A# set ae1 aggregated-ether-options link-speed 10g
```

6. Specify the interface ID of the uplinks to be included in LAG **ae0**:

```
[edit interfaces]
user@Host-A# set ge-0/1/0 ether-options 802.ad ae0
user@Host-A# set ge-1/1/0 ether-options 802.ad ae0
```

7. Specify the interface ID of the uplinks to be included in LAG **ae1**:

```
[edit interfaces]
user@Host-A# set ge-0/1/1 ether-options 802.ad ae1
user@Host-A# set ge1/1/1 ether-options 802.ad ae1
```

8. Specify that LAG **ae0** belongs to the subnet for the employee broadcast domain:

```
[edit interfaces]
user@Host-A# set ae0 unit 0 family inet address 192.0.2.0/25
```

9. Specify that LAG **ae1** belongs to the subnet for the guest broadcast domain:

```
[edit interfaces]
user@Host-A# set ae1 unit 1 family inet address 192.0.2.128/25
```

Results Display the results of the configuration:

```
[edit]
chassis {
  aggregated-devices {
    ethernet {
      device-count 2;
    }
  }
}
interfaces {
  ae0 {
    aggregated-ether-options {
      link-speed 10g;
      minimum-links 2;
    }
    unit 0 {
      family inet {
        address 192.0.2.0/25;
```

```
        }
      }
    }
    ae1 {
      aggregated-ether-options {
        link-speed 10g;
        minimum-links 2;
      }
      unit 0 {
        family inet {
          address 192.0.2.128/25;
        }
      }
    ge–0/1/0 {
      ether-options {
        802.ad ae0;
      }
    }
    ge–1/1/0 {
      ether-options {
        802.ad ae0;
      }
    }
    ge–0/1/1 {
      ether-options {
        802.ad ae1;
      }
    }
    gxe–1/1/1 {
      ether-options {
        802.ad ae1;
      }
    }
  }
```

Verification

To verify that switching is operational and two LAGs have been created, perform these tasks:

- Verifying That LAG ae0 Has Been Created on page 745
- Verifying That LAG ae1 Has Been Created on page 746

Verifying That LAG ae0 Has Been Created

Purpose Verify that LAG **ae0** has been created on the switch.

Action **show interfaces ae0 terse**

```
Interface               Admin      Link Proto      Local              Remote

ae0                      up       up
ae0.0                    up       up     inet     10.10.10.2/24
```

Meaning The output confirms that the **ae0** link is up and shows the **family** and IP address assigned to this link.

Verifying That LAG ae1 Has Been Created

Purpose Verify that LAG **ae1** has been created on the switch

Action **show interfaces ae1 terse**

```
Interface               Admin Link Proto    Local                 Remote
ae1                     up    down
ae1.0                   up    down inet
```

Meaning The output shows that the **ae1** link is down.

Troubleshooting

Troubleshooting a LAG That Is Down

Problem The **show interfaces terse** command shows that the LAG is **down:**

Solution Check the following:

- Verify that there is no configuration mismatch.
- Verify that all member ports are up.
- Verify that a LAG is part of family ethernet switching (Layer 2 LAG) or family inet (Layer 3 LAG).
- Verify that the LAG member is connected to the correct LAG at the other end.
- Verify that the LAG members belong to the same switch (or the same Virtual Chassis).

Related Documentation

- Example: Configuring a Virtual Chassis with a Master and Backup in a Single Wiring Closet on page 717
- Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 746
- Example: Connecting an Access Switch to a Distribution Switch on page 1078.
- Virtual Chassis Cabling Configuration Examples for J-EX4200 Switches
- Installing an Uplink Module in a J-EX4200 Switch

Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch

J-EX Series switches allow you to combine multiple Ethernet links into one logical interface for higher bandwidth and redundancy. The ports that are combined in this manner are referred to as a link aggregation group (LAG) or bundle. The number of Ethernet links you can combine into a LAG depends on your J-EX Series switch model. See “Understanding Aggregated Ethernet Interfaces and LACP” on page 867 for more

information. J-EX Series switches allow you to further enhance these links by configuring Link Aggregation Control Protocol (LACP).

This example describes how to overlay LACP on the LAG configurations that were created in "Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch" on page 740:

Requirements

This example uses the following software and hardware components:

- Two J-EX4200-48T switches
- Two J-EX4200-24F switches
- Four J-EX Series uplink modules

Before you configure LACP, be sure you have:

- Set up the Virtual Chassis switches. See "Example: Configuring a Virtual Chassis with a Master and Backup in a Single Wiring Closet" on page 717.
- Configured the uplink ports on the switches as trunk ports. See "Configuring Gigabit Ethernet Interfaces (CLI Procedure)" on page 919.
- Configured the LAGs. See "Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch" on page 740

Overview and Topology

This example assumes that you are already familiar with the Example: Configuring Aggregated Ethernet High-Speed Uplinks between Virtual Chassis Access Switch and Virtual Chassis Distribution Switch. The topology in this example is exactly the same as the topology in that other example. This example shows how to use LACP to enhance the LAG functionality.

LACP exchanges are made between *actors* (the transmitting link) and *partners* (the receiving link). The LACP *mode* can be either active or passive.

NOTE: If the actor and partner are both in passive mode, they do not exchange LACP packets, which results in the aggregated Ethernet links not coming up. By default, LACP is in passive mode. To initiate transmission of LACP packets and responses to LACP packets, you must enable LACP in active mode.

By default, the actor and partner send LACP packets every second. You can configure the interval at which the interfaces send LACP packets by including the periodic statement at the **[edit interfaces *interface-name* aggregated-ether-options lacp]** hierarchy level.

The interval can be fast (every second) or slow (every 30 seconds).

Configuring LACP for the LAGs on the Virtual Chassis Access Switch

To configure LACP for the access switch LAGs, perform these tasks:

CLI Quick Configuration

To quickly configure LACP for the access switch LAGs, copy the following commands and paste them into the switch terminal window:

```
[edit]
set interfaces ae0 aggregated-ether-options lacp active periodic fast
set interfaces ae1 aggregated-ether-options lacp active periodic fast
```

Step-by-Step Procedure

To configure LACP for Host-A LAGs **ae0** and **ae1**:

1. Specify the aggregated Ethernet options for both bundles:

```
[edit interfaces]
user@Host-A#set ae0 aggregated-ether-options lacp active periodic fast
user@Host-A#set ae1 aggregated-ether-options lacp active periodic fast
```

Results

Display the results of the configuration:

```
[edit interfaces]
user@Host-A# show
ae0 {
  aggregated-ether-options {
    lacp {
      active;
      periodic fast;
    }
  }
}
ae1 {
  aggregated-ether-options {
    lacp {
      active;
      periodic fast;
    }
  }
}
```

Configuring LACP for the LAGs on the Virtual Chassis Distribution Switch

To configure LACP for the two uplink LAGs from the Virtual Chassis access switch to the Virtual Chassis distribution switch, perform these tasks:

CLI Quick Configuration

To quickly configure LACP for the distribution switch LAGs, copy the following commands and paste them into the switch terminal window:

```
[edit interfaces]
set ae0 aggregated-ether-options lacp passive periodic fast
set ae1 aggregated-ether-options lacp passive periodic fast
```

Step-by-Step Procedure

To configure LACP for Host D LAGs **ae0** and **ae1**:

1. Specify the aggregated Ethernet options for both bundles:

```
[edit interfaces]
user@Host-D#set ae0 aggregated-ether-options lacp passive periodic fast
user@Host-D#set ae1 aggregated-ether-options lacp passive periodic fast
```

Results

Display the results of the configuration:

```
[edit interfaces]
user@Host-D# show
ae0 {
  aggregated-ether-options {
    lacp {
      passive;
      periodic fast;
    }
  }
}
ae1 {
  aggregated-ether-options {
    lacp {
      passive
      periodic fast;
    }
  }
}
```

Verification

To verify that LACP packets are being exchanged, perform these tasks:

- Verifying the LACP Settings on page 749
- Verifying That the LACP Packets Are Being Exchanged on page 750

Verifying the LACP Settings

Purpose Verify that LACP has been set up correctly.

Action Use the **show lacp interfaces *interface-name*** command to check that LACP has been enabled as active on one end.

```
user@Host-A> show lacp interfaces xe-0/1/0

Aggregated interface: ae0

    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity

      ge-0/1/0      Actor    No   Yes    No   No   No   Yes     Fast    Active

      ge-0/1/0    Partner    No   Yes    No   No   No   Yes     Fast   Passive

    LACP protocol:   Receive State    Transmit State          Mux State

      ge-0/1/0          Defaulted     Fast periodic           Detached
```

Meaning The output indicates that LACP has been set up correctly and is active at one end.

Verifying That the LACP Packets Are Being Exchanged

Purpose Verify that LACP packets are being exchanged.

Action Use the **show interfaces ae*x* statistics** command to display LACP information.

```
user@Host-A> show interfaces ae0 statistics

Physical interface: ae0, Enabled, Physical link is Down
  Interface index: 153, SNMP ifIndex: 30
  Link-level type: Ethernet, MTU: 1514, Speed: Unspecified, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1,
  Minimum bandwidth needed: 0
  Device flags   : Present Running
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Current address: 02:19:e2:50:45:e0, Hardware address: 02:19:e2:50:45:e0
  Last flapped   : Never
  Statistics last cleared: Never
    Input packets : 0
    Output packets: 0
  Input errors: 0, Output errors: 0

  Logical interface ae0.0 (Index 71) (SNMP ifIndex 34)
    Flags: Hardware-Down Device-Down SNMP-Traps Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :             0          0             0            0
        Output:             0          0             0            0
    Protocol inet
      Flags: None
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255
```

Meaning The output here shows that the link is down and that no PDUs are being exchanged.

Troubleshooting

These are some tips for troubleshooting:

Troubleshooting a Nonworking LACP Link

Problem The LACP link is not working.

Solution Check the following:

- Remove the LACP configuration and verify whether the static LAG is up.
- Verify that LACP is configured at both ends.

- Verify that LACP is not passive at both ends.
- Verify whether LACP protocol data units (PDUs) are being exchanged by running the **monitor traffic-interface** ***lag-member*** **detail** command.

Related Documentation

- Example: Connecting an Access Switch to a Distribution Switch on page 1078
- Virtual Chassis Cabling Configuration Examples for J-EX4200 Switches
- Installing an Uplink Module in a J-EX4200 Switch

Example: Configuring a Virtual Chassis Using a Preprovisioned Configuration File

You can deterministically control both the role and the member ID assigned to each member switch in a Virtual Chassis configuration by creating a preprovisioned configuration file.

A preprovisioned configuration file links the serial number of each J-EX4200 switch in the configuration to a specified member ID and role. The serial number must be specified in the configuration file for the member to be recognized as part of the Virtual Chassis configuration.

You must select two members that you want to make eligible for election as master of the Virtual Chassis configuration. When you list these two members in the preprovisioned configuration file, you designate both members as **routing-engine**. One will function as the master of the Virtual Chassis configuration and the other will function as the backup.

You designate additional members, which are not eligible for election as master, has having the **linecard** role in the preprovisioned configuration file.

NOTE: When you use a preprovisioned configuration, you cannot modify the mastership priority or member ID of member switches through the user interfaces.

NOTE: After you have created a preprovisioned Virtual Chassis configuration, you can use the autoprovisioning feature to add member switches to that configuration. See “Adding a New Switch to an Existing Virtual Chassis Configuration (CLI Procedure)” on page 786.

This example describes how to configure a Virtual Chassis across multiple wiring closets using a preprovisioned configuration file:

Requirements

This example uses the following hardware and software components:

- Five J-EX4200-48T switches
- Five J-EX4200-24T switches
- Four uplink modules

Before you create the preprovisioned configuration of the Virtual Chassis and interconnect the members across the wiring closets, be sure you have:

1. Made a list of the serial numbers of all the switches to be connected as a Virtual Chassis configuration.
2. Noted the desired role (**routing-engine** or **linecard**) of each switch. If you configure the member with a **routing-engine** role, it is eligible to function as a master or backup. If you configure the member with a **linecard** role , it is not eligible to become a master or backup.
3. Installed an uplink module in each of the member switches that will be interconnected across wiring closets. See Installing an Uplink Module in a J-EX4200 Switch.
4. Interconnected the member switches within each wiring closet using the dedicated VCPs on the rear panel of switches. See Connecting a Virtual Chassis Cable to a J-EX4200 Switch.
5. Powered on the switch that you plan to use as the master switch (SWA-0).
6. Run the EZSetup program on SWA-0, specifying the identification parameters. See “Connecting and Configuring a J-EX Series Switch (CLI Procedure)” on page 161 for details.

   SWA-0 is going to be configured in the example to function as the master of the Virtual Chassis configuration. Thus, the properties that you specify for SWA-0 will apply to the entire Virtual Chassis configuration, including all the member switches that you specify in the preprovisioned configuration file.
7. Configured SWA-0 with the virtual management Ethernet (VME) interface for out-of-band management of the Virtual Chassis configuration, if desired.

   ```
   [edit]
   user@SWA-0# set interfaces vme unit 0 family inet address /ip-address/mask/
   ```

Overview and Topology

In this example, five J-EX4200 switches (SWA-0 through SWA-4) are interconnected with their dedicated VCPs in wiring closet A and five J-EX4200 switches (SWA-5 through SWA-9) are interconnected with their dedicated VCPs in wiring closet B.

SWA-0 (in wiring closet A) is going to be the master of the Virtual Chassis configuration. This example shows how to create a preprovisioned configuration file on SWA-0 for all member switches that will be interconnected in the Virtual Chassis configuration. The preprovisioned configuration file includes member IDs for the members in wiring closet A and for the members in wiring closet B.

SWA-5 (in wiring closet B) is going to be the backup of the Virtual Chassis configuration. Both SWA-0 and SWA-5 are specified in the preprovisioned configuration file with the role of **routing-engine**. All other members are specified with the role of **linecard**.

If all member switches could be interconnected with their dedicated VCPs, you could simply power on the switches after saving and committing the preprovisioned configuration file. The master detects the connection of the members through the

dedicated VCPs and applies the parameters specified in the preprovisioned configuration file.

However, the Virtual Chassis cables that interconnect the VCPs of member switches within a single wiring closet are not long enough to connect member switches across wiring closets. Instead, you can configure the uplink module ports and the SFP networks ports on J-EX4200-24F switches as VCPs to interconnect the member switches in wiring closet A to the member switch in wiring closet B. For redundancy, this example connects uplink VCPs from two member switches in wiring closet A (SWA–0 and SWA–2) to two member switches (SWA-5 and SWA-7) in wiring closet B.

NOTE: You can use interfaces on SFP and SFP+ uplink modules and the SFP network ports on J-EX4200-24F switches as VCPs. When an uplink module port or SFP network port is set as a VCP, it cannot be used for any other purpose. The SFP uplink module has four 1-Gbps ports; the SFP+ uplink module has four 1-Gbps or two 10-Gbps ports. The uplink module ports that are not set as VCPs can be configured as trunk ports to connect to a distribution switch.

Because this particular preprovisioned configuration is for a Virtual Chassis that is interconnected across wiring closets, we will bring up the Virtual Chassis configuration in stages. First, we power on SWA-0 (without powering on any other switches) and create the preprovisioned configuration file. Then we power on the remaining switches in wiring closet A. If we check the status of the Virtual Chassis configuration at this point by using the **show virtual-chassis status** command, it will display only **member 0** through **member 4**. The members that have not yet been interconnected will not be listed.

Next power on SWA-5 without powering on the remaining switches (SWA-6 through SWA-9) in wiring closet B. Bring up SWA-5 as a standalone switch and set one of its uplinks as a VCP prior to interconnecting it with the Virtual Chassis configuration in wiring closet A. Without this setting, SWA-5 cannot be detected as a member switch by the master of the Virtual Chassis configuration.

You can set the uplink VCP of SWA–5 without running the EZSetup program by directly connecting to the console port. If you wish, you can run the EZSetup program and specify identification parameters. When you interconnect SWA-5 with the master of the Virtual Chassis configuration, the master overwrites any conflicting parameters.

After setting the VCP in SWA-5, connect this VCP with the VCP of SWA-0 in wiring closet A. SWA-5 (serial number pqr678) is specified as a **routing-engine** in the preprovisioned configuration file.

This example uses SWA-5 as the backup of the Virtual Chassis configuration. If a problem occurred in wiring closet A, SWA-5 would take control of the Virtual Chassis configuration and maintain the network connections. Specify both SWA-0 and SWA-5 as **routing-engine**. Because SWA-0 is powered on prior to SWA-5, it has additional prioritization properties that cause it to be elected as master of the Virtual Chassis configuration.

After being physically interconnected with SWA-0, SWA-5 reboots and comes up as **member 5** and as the backup of the Virtual Chassis configuration.

Power on the remaining switches (SWA-6 through SWA-9) in wiring closet B. The master can now detect that all members are present. Finally, for redundancy, configure an additional VCP on SWA-7 through the master.

The topology for this example consists of:

- Three J-EX4200-48T switches (SWA-0 , SWA-2, and SWA-4) in wiring closet A.
- Two J-EX4200-48T switches (SWA-5 and SWA-9) in wiring closet B.
- Two J-EX4200-24T switches (SWA-1 and SWA-3) in wiring closet A.
- Three J-EX4200-24T switches (SWA-6, SWA-7, and SWA-8) in wiring closet B.
- Four uplink modules. Two are installed in wiring closet A and two are installed in wiring closet B.

Table 106 on page 755 shows the Virtual Chassis configuration settings for a preprovisioned Virtual Chassis composed of member switches in different wiring closets.

Table 106: Components of a Preprovisioned Virtual Chassis Interconnected Across Multiple Wiring Closets

Switch	Serial number	Member ID	Role	Uplink Module Ports	Hardware	Location
SWA-0	abc123	0	routing-engine	**ge-0/1/0**	J-EX4200-48T and uplink module	Wiring closet A
SWA-1	def456	1	linecard		J-EX4200-24T	Wiring closet A
SWA-2	ghi789	2	linecard	**ge-2/1/0**	J-EX4200-48T and uplink module	Wiring closet A
SWA-3	jkl012	3	linecard		J-EX4200-24T	Wiring closet A
SWA-4	mno345	4	linecard		J-EX4200-48T	Wiring closet A
SWA-5	pqr678	5	routing-engine	**ge-0/1/0** NOTE: The member ID of SWA-5 is **0** at the time that its uplink module port is configured as a VCP.	J-EX4200-48T and uplink module	Wiring closet B
SWA-6	stu901	6	linecard		J-EX4200-24T	Wiring closet B

Table 106: Components of a Preprovisioned Virtual Chassis Interconnected Across Multiple Wiring Closets *(continued)*

Switch	Serial number	Member ID	Role	Uplink Module Ports	Hardware	Location
SWA-7	vwx234	7	linecard	**ge-7/1/0**	J-EX4200-24T and uplink module	Wiring closet B
SWA-8	yza567	8	linecard		J-EX4200-24T	Wiring closet B
SWA-9	bcd890	9	linecard		J-EX4200-48T	Wiring closet B

Figure 19 on page 757 shows the different types of interconnections used for this Virtual Chassis configuration. The rear view shows that the member switches within each wiring closet are interconnected to each other using the dedicated VCPs. The front view shows that the uplink module ports that have been set as VCPs and interconnected across the wiring closets. The uplink module ports that are not set as VCPs can be configured as trunk ports to connect to a distribution switch.

NOTE: The interconnections shown in Figure 19 on page 757 are the same as they would be for a configuration that was not preprovisioned across wiring closets.

Figure 19: Maximum Size Virtual Chassis Interconnected Across Wiring Closets

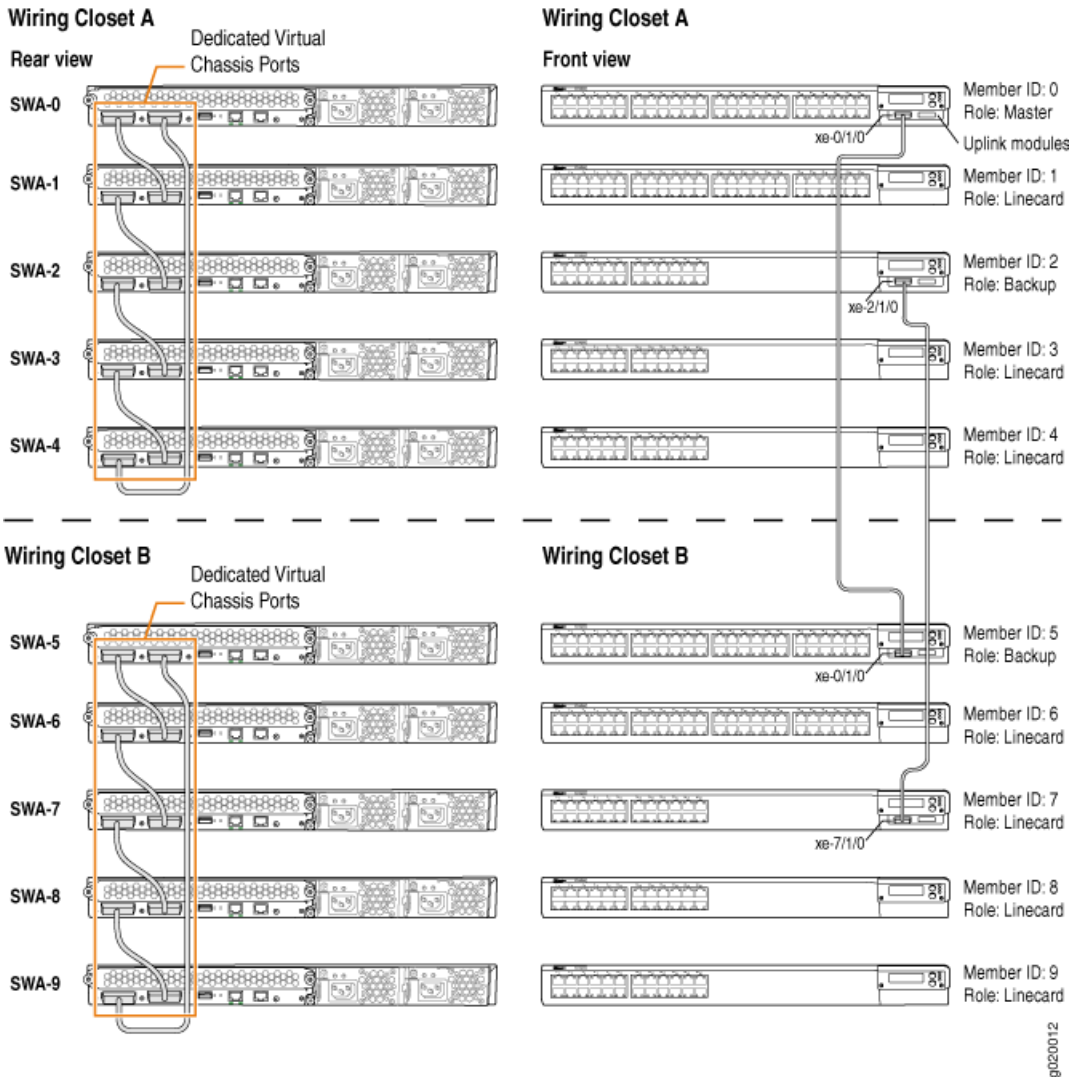

Configuration

To configure the Virtual Chassis across multiple wiring closets using a preprovisioned configuration:

NOTE: We recommend that you use the commit synchronize command to save any configuration changes that you make to a multimember Virtual Chassis configuration.

Step-by-Step Procedure

To create a preprovisioned configuration for the Virtual Chassis:

1. Specify the preprovisioned configuration mode:

```
[edit virtual-chassis]
user@SWA-0# set preprovisioned
```

2. Specify all the members that will be included in the Virtual Chassis configuration, listing each switch's serial number with the desired member ID and the desired role:

```
[edit virtual-chassis]
user@SWA-0# set member 0 serial-number abc123 role routing-engine
user@SWA-0# set member 1 serial-number def456 role linecard
user@SWA-0# set member 2 serial-number ghi789 role linecard
user@SWA-0# set member 3 serial-number jkl012 role linecard
user@SWA-0# set member 4 serial-number mno345 role linecard
user@SWA-0# set member 5 serial-number pqr678 role routing-engine
user@SWA-0# set member 6 serial-number stu901 role linecard
user@SWA-0# set member 7 serial-number vwx234 role linecard
user@SWA-0# set member 8 serial-number yza567 role linecard
user@SWA-0# set member 9 serial-number bcd890 role linecard
```

3. Power on the member switches in wiring closet A.

4. Prepare the members in wiring closet A for interconnecting with the member switches in wiring closet B by setting uplink VCPs for member 0 and member 2:

```
user@SWA-0> request virtual-chassis vc-port set pic-slot 1 port 0
user@SWA-2> request virtual-chassis vc-port set pic-slot 1 port 0 member 2
```

NOTE:

- For redundancy, this example sets an uplink VCP in both SWA-0 and SWA-2.
- This example omits the specification of the member 0 in setting the uplink for SWA-0. The command applies by default to the switch where it is executed.

5. Power on SWA-5 and connect to it. This switch comes up as member ID **0** and functions as master of itself. Although SWA-5 is listed in the preprovisioned configuration file, it is not a present member of the Virtual Chassis configuration that has been powered on thus far. In order for the master to detect SWA-5 as a connected member, you must first set an uplink VCP on SWA-5 and interconnect that VCP with the uplink VCP of SWA-0.

6. Set the first uplink of SWA-5 to function as a VCP. Because SWA-5 has been powered on as a separate switch and is still operating independently at this point, its member ID is **0**.

```
user@SWA-5> request virtual-chassis vc-port set pic-slot 1 port 0
```

NOTE: This example omits the specification of the member 0 in configuring the uplink for SWA-5 (at this point the member ID of SWA-5 is still 0). The command applies by default to the switch where it is executed.

7. Power off SWA-5 and connect the fiber cable from SWA-5 uplink VCP **ge-0/1/0** to the uplink VCP **ge-0/1/0** on SWA-0.
8. Power on SWA-5.
9. Now that SWA-5 has been brought up as **member 5** of the Virtual Chassis configuration, power on the remaining switches (SWA-6 through SWA-9) in wiring closet B. They are interconnected with SWA-5 using the dedicated VCPs on the rear panel and are therefore detected by the master as interconnected members. If you check the status of the Virtual Chassis configuration at this point, all the members that were specified in the preprovisioned configuration file should be displayed as present. Additional configuration for member switches can now be done through the master switch.
10. Set one uplink module port of SWA-7 to function as a VCP:

```
user@SWA-0> request virtual-chassis vc-port set pic-slot 1 port 0
member 7
```

Results Display the results of the configuration on SWA-0:

```
[edit]
user@SWA-0# show
virtual-chassis {
  member 0 {
    role routing-engine;
    serial-number abc123;
  }
  member 1 {
    role linecard;
    serial-number def456;
  }
  member 2 {
    role linecard;
    serial-number ghi789;
  }
  member 3 {
    role linecard;
    serial-number jkl012;
  }
  member 4 {
    role linecard;
    serial-number mno345;
  }
  member 5 {
    role routing-engine;
    serial-number pqr678;
  }
```

```
    member 6 {
      role linecard;
      serial-number stu901;
    }
    member 7 {
      role linecard;
      serial-number vwx234;
    }
    member 8 {
      role linecard;
      serial-number yza567;
    }
    member 9 {
      role linecard;
      serial-number bcd890;
    }
    preprovisioned;
  }
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying the Member IDs and Roles of the Member Switches on page 760
- Verifying That the Dedicated VCPs and Uplink VCPs Are Operational on page 761

Verifying the Member IDs and Roles of the Member Switches

Purpose Verify that the member IDs and roles are all set as expected.

Action Display the members of the Virtual Chassis configuration:

```
user@SWA-0> show virtual-chassis status
Preprovisioned Virtual Chassis
Virtual Chassis ID: 0000.e255.0000
                                                Mastership          Neighbor List
Member ID  Status   Serial No    Model         Priority  Role      ID  Interface

0 (FPC 0)  Prsnt    abc123       ex4200-48t      129  Master*    1  vcp-0
                                                                  4  vcp-1
                                                                  5  1/0

1 (FPC 1)  Prsnt    def456       ex4200-24t        0  Linecard   2  vcp-0
                                                                 0  vcp–1


2 (FPC 2)  Prsnt    ghi789       ex4200-48t        0  Linecard   3  vcp-0
                                                                  1  vcp-1
                                                                  7  1/0

3 (FPC 3)  Prsnt    jkl012       ex4200-24t        0  Linecard   4  vcp-0
                                                                  2  vcp-1

4 (FPC 4)  Prsnt    mno345       ex4200-48t        0  Linecard   0  vcp-0
                                                                  3  vcp-1

5  FPC 5)  Prsnt    pqr678       ex4200-48t      129  Backup     6  vcp-0
                                                                  9  vcp-1
```

```
                                                                    0  1/0

6 (FPC 6)  Prsnt    stu901       ex4200-24t        0  Linecard   7  vcp-0
                                                                    5  vcp-1

7 (FPC 7)  Prsnt    vwx234       ex4200-24t        0  Linecard   8  vcp-0
                                                                    6  vcp-1
                                                                    2  1/0

8 (FPC 8)  Prsnt    yza567       ex4200-24t        0  Linecard   9  vcp-0
                                                                    7  vcp-1

9 (FPC 9)  Prsnt    bc7890       ex4200-48t        0  Linecard   5  vcp-0
                                                                    8  vcp-1
```

Meaning The output shows that all members listed in the preprovisioned configuration file are connected to the Virtual Chassis configuration. It confirms that SWA-0 (member **0**) is functioning as the master of the Virtual Chassis configuration, which was the intention of the configuration procedure. The other switch configured with the **routing-engine** role (SWA-5) is functioning as the backup. The **Neighbor List** displays the interconnections of the member VCPs.

Verifying That the Dedicated VCPs and Uplink VCPs Are Operational

Purpose Verify that the dedicated VCPs interconnecting the member switches within each wiring closet and the uplink module VCPs interconnecting the member switches across wiring closets are operational.

Action Display the Virtual Chassis interfaces:

```
user@SWA-0> show virtual-chassis vc-port all-members
fpc0:
--------------------------------------------------------------------------
Interface        Type            Status    Speed        Neighbor
or                                         (mbps)     ID   Interface
PIC / Port
vcp-0            Dedicated       Up
vcp-1            Dedicated       Up
1/0              Configured      Up

fpc1:
--------------------------------------------------------------------------
Interface        Type            Status    Speed        Neighbor
or                                         (mbps)     ID   Interface
PIC / Port
vcp-0            Dedicated       Up
vcp-1            Dedicated       Up

fpc2:
--------------------------------------------------------------------------
Interface        Type            Status    Speed        Neighbor
or                                         (mbps)     ID   Interface
PIC / Port
vcp-0            Dedicated       Up
vcp-1            Dedicated       Up
1/0              Configured      Up
```

```
fpc3:
--------------------------------------------------------------------------
Interface        Type            Status    Speed        Neighbor
or                                         (mbps)     ID   Interface
PIC / Port
vcp-0            Dedicated       Up
vcp-1            Dedicated       Up

fpc4:
--------------------------------------------------------------------------
Interface        Type            Status    Speed        Neighbor
or                                         (mbps)     ID   Interface
PIC / Port
vcp-0            Dedicated       Up
vcp-1            Dedicated       Up

fpc5:
--------------------------------------------------------------------------
Interface        Type            Status    Speed        Neighbor
or                                         (mbps)     ID   Interface
PIC / Port
vcp-0            Dedicated       Up
vcp-1            Dedicated       Up
1/0              Configured      Up

fpc6:
--------------------------------------------------------------------------
Interface        Type            Status    Speed        Neighbor
or                                         (mbps)     ID   Interface
PIC / Port
vcp-0            Dedicated       Up
vcp-1            Dedicated       Up

fpc7:
--------------------------------------------------------------------------
Interface        Type            Status    Speed        Neighbor
or                                         (mbps)     ID   Interface
PIC / Port
vcp-0            Dedicated       Up
vcp-1            Dedicated       Up
1/0              Configured      Up

fpc8:
--------------------------------------------------------------------------
Interface        Type            Status    Speed        Neighbor
or                                         (mbps)     ID   Interface
PIC / Port
vcp-0            Dedicated       Up
vcp-1            Dedicated       Up


fpc9:
--------------------------------------------------------------------------
Interface        Type            Status    Speed        Neighbor
or                                         (mbps)     ID   Interface
PIC / Port
vcp-0            Dedicated       Up
```

```
vcp-1           Dedicated       Up
```

Meaning The dedicated VCPs interconnecting the member switches within wiring closets are displayed as **vcp-0** and **vcp-1**. The uplink module VCPs interconnecting member switches (members 0, 2, 5, and 7) across wiring closets are displayed as **1/0** and **1/1** and identified as **Configured**.

Troubleshooting

To troubleshoot a preprovisioned Virtual Chassis configuration that is interconnected across wiring closets, perform these tasks:

Troubleshooting Nonoperational VCPs

Problem A VCP shows a status of **down**.

Solution Check the cable to make sure that it is properly and securely connected to the ports.

Related Documentation

- Example: Configuring a Virtual Chassis with a Master and Backup in a Single Wiring Closet on page 717
- Example: Configuring a Virtual Chassis Interconnected Across Multiple Wiring Closets on page 733
- Configuring a Virtual Chassis (CLI Procedure) on page 781
- Configuring a Virtual Chassis (J-Web Procedure) on page 784

Example: Configuring Fast Failover on Uplink Module VCPs to Reroute Traffic When a Virtual Chassis Member Switch or Intermember Link Fails

The Virtual Chassis fast failover feature is a hardware-assisted failover mechanism that automatically reroutes traffic and reduces traffic loss in the event of a link or switch failure. If a link between two members fails, traffic flow between those members must be rerouted quickly so that there is minimal traffic loss.

Fast failover is enabled by default on all dedicated Virtual Chassis ports (VCPs).

This example describes how to configure fast failover on uplink module VCPs in a Virtual Chassis configuration:

Requirements

This example uses the following hardware and software components:

- Six J-EX4200-24T switches
- Four SFP uplink modules

Before you begin configuring fast failover, be sure you have:

1. Mounted the switches. See Mounting a J-EX4200 Switch on Two Posts in a Rack or Cabinet, Mounting a J-EX4200 Switch on a Desk or Other Level Surface, or Mounting a J-EX4200 Switch on a Wall.
2. Cabled the switches in a multiple-ring topology to create the Virtual Chassis configuration. See Connecting a Virtual Chassis Cable to a J-EX4200 Switch and “Example: Configuring a Virtual Chassis Interconnected Across Multiple Wiring Closets” on page 733. See Figure 20 on page 765 for an illustration of a multiple-ring topology.

Overview and Topology

In a Virtual Chassis configuration, fast failover automatically reroutes traffic and reduces traffic loss in the event of a link failure or a member switch failure. By default, fast failover is enabled on all dedicated Virtual Chassis ports (VCPs). If you configure uplink module ports as VCPs, you must manually configure fast failover on these ports.

For fast failover to be effective, the Virtual Chassis members must be configured in a ring topology. The ring topology can be formed by using either dedicated Virtual Chassis ports (VCPs) or user-configured uplink module VCPs. Fast failover is supported only in a ring topology that uses identical port types, for example, either a topology that uses all dedicated VCPs or one that uses all uplink module VCPs. Fast failover is not supported in a ring topology that includes both dedicated VCPs and uplink module VCPs. Fast failover is supported, however, in a Virtual Chassis configuration that consists of multiple rings.

Figure 20 on page 765 shows an example of a multiple-ring topology.

Figure 20: Traffic Redirected by Fast Failover After VCP Link Failures in a Topology with Multiple Rings

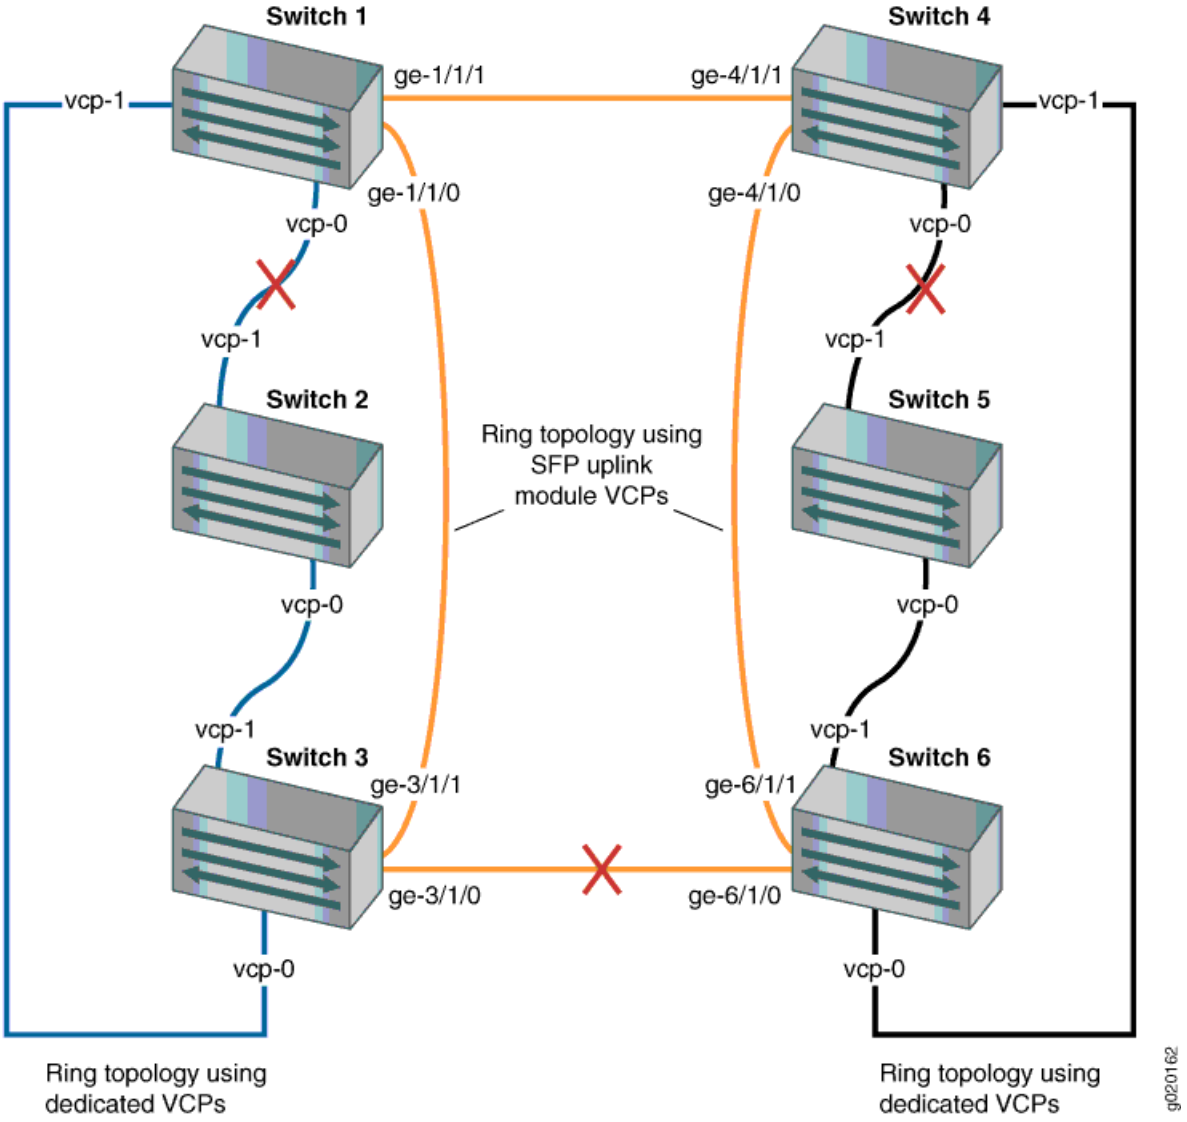

This example shows how to enable fast failover on uplink module VCPs.

The topology for this example consists of six switches:

- Six J-EX4200-24T switches, four of which have an SFP uplink module installed (switches 1, 3, 4, and 6)

Configuration

To configure the fast failover feature on uplink module VCPs:

CLI Quick Configuration

To configure fast failover on all SFP uplink module VCPs, copy the following command and paste it into the terminal window on switch 1:

```
[edit]
set virtual-chassis fast-failover ge
```

Step-by-Step Procedure

To configure fast failover on SFP uplink module VCPs:

1. Enable fast failover on all SFP uplink module VCPs in the Virtual Chassis configuration:

```
[edit]
user@switch1# set virtual-chassis fast-failover ge
```

NOTE: We recommend that you use the **commit synchronize** command to save any configuration changes that you make to a multimember Virtual Chassis.

Results

Check the results of the configuration:

```
[edit virtual-chassis]
user@switch1# show
fast-failover {
  ge;
}
```

Verification

To confirm that fast failover is enabled on SFP uplink module VCPs in the Virtual Chassis configuration, perform these tasks:

- Verifying That Fast Failover Is Enabled on page 766

Verifying That Fast Failover Is Enabled

Purpose

Verify that fast failover has been enabled in a Virtual Chassis configuration.

Action

1. Issue the **show virtual-chassis fast-failover** command.
2. Check to see that fast failover is enabled.

user@switch1> **show virtual-chassis fast-failover**

```
Fast failover on dedicated VCP ports: Enabled
Fast failover on XE uplink VCP ports: Disabled
Fast failover on GE uplink VCP ports: Enabled
```

Meaning

Fast failover is enabled on all dedicated VCPs and SFP uplink module VCPs in the Virtual Chassis configuration.

Related Documentation

- Configuring Fast Failover in a Virtual Chassis Configuration on page 798
- Disabling Fast Failover in a Virtual Chassis Configuration on page 799

Example: Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge

You can explicitly assign a Virtual Chassis ID so that, when two Virtual Chassis configurations merge, the ID that you assigned takes precedence over the automatically assigned Virtual Chassis IDs and becomes the ID of the newly merged Virtual Chassis configuration.

This example describes how to assign the Virtual Chassis ID in a Virtual Chassis configuration:

Requirements

This example uses the following hardware and software components:

- Two J-EX4200-48T switches
- Two J-EX4200-24T switches

Before you begin, be sure you have:

1. Installed the switches. See Mounting a J-EX4200 Switch on Two Posts in a Rack or Cabinet, Mounting a J-EX4200 Switch on a Desk or Other Level Surface, or Mounting a J-EX4200 Switch on a Wall.
2. Cabled the switches to create the Virtual Chassis configuration. See Connecting a Virtual Chassis Cable to a J-EX4200 Switch.

Overview and Topology

Every Virtual Chassis configuration has a unique ID that is automatically assigned when the Virtual Chassis configuration is formed. You can also configure a Virtual Chassis ID using the **set virtual-chassis id** command. When two Virtual Chassis merge, the Virtual Chassis ID that you assigned takes precedence over the automatically assigned Virtual Chassis IDs and becomes the ID for the newly merged Virtual Chassis configuration.

The topology for this example consists of four switches:

- Two J-EX4200-24T switches
- Two J-EX4200-48T switches

The switches are connected as a four-member Virtual Chassis configuration and are identified as switch-A, switch-B, switch-C, and switch-D. The master is switch-A.

Configuration

Assign the Virtual Chassis ID in a Virtual Chassis configuration:

CLI Quick Configuration

To assign a Virtual Chassis ID so that, when two Virtual Chassis configurations merge, the ID that you assigned takes precedence over the automatically assigned Virtual Chassis IDs and becomes the ID of the newly merged Virtual Chassis configuration, copy the following command and paste it into the terminal window:

```
[edit]
set virtual-chassis id 9622.6ac8.5345
```

Step-by-Step Procedure

To assign the Virtual Chassis ID in a Virtual Chassis configuration:

1. Assign the Virtual Chassis ID:

```
[edit]
user@switch-A# set virtual-chassis id 9622.6ac8.5345
```

NOTE: We recommend that you use the **commit synchronize** command to save any configuration changes that you make to a multimember Virtual Chassis configuration.

Verification

To verify that the Virtual Chassis ID has been assigned as you intended, perform these tasks:

- Verifying That the Virtual Chassis ID Is Assigned on page 768

Verifying That the Virtual Chassis ID Is Assigned

Purpose Verify that the Virtual Chassis ID has been assigned in a Virtual Chassis configuration.

Action

1. Issue the **show configuration virtual-chassis id** command.
2. Check to see that the Virtual Chassis ID number is listed.

```
user@switch-A> show configuration virtual-chassis id
id 9622.6ac8.5345;
```

Meaning The Virtual Chassis ID has been assigned as 9622.6ac8.5345.

Related Documentation

- Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge (CLI Procedure) on page 800
- Configuring a Virtual Chassis (CLI Procedure) on page 781
- Configuring a Virtual Chassis (J-Web Procedure) on page 784

Example: Configuring Link Aggregation Groups Using Uplink Virtual Chassis Ports

You can form link aggregation groups (LAGs) between Virtual Chassis member switches in different wiring closets using uplink Virtual Chassis ports (VCPs) and, on J-EX4200-24F switches, network VCPs. LAGs balance traffic across the member links, increase the uplink bandwidth, and provide increased availability. To form LAGs using uplink or network VCPs, you configure the uplink module interfaces or network interfaces on the member switches as VCPs and connect the VCPs using fiber-optic cables. For the LAGs to form, the uplink or network VCPs on each member switch that will form a LAG must operate at the same link speed and you must interconnect at least two uplink or network VCPs on each of those members. You can connect uplink or network VCPs operating at different link speeds, but they will not form a LAG.

NOTE: The LAGs formed by VCPs are different from LAGs formed by Virtual Chassis network interfaces. For more information on LAGs formed by network interfaces, see “Understanding Virtual Chassis Configurations and Link Aggregation” on page 702.

This example shows how to configure uplink module interfaces and network interfaces as VCPs on multiple member switches of a Virtual Chassis configuration and then connect them to form LAGs:

Requirements

This example uses the following hardware and software components:

- Five J-EX4200 switches, one of which is a J-EX4200-24F model
- Four uplink modules

Before you configure the uplink module interfaces and network interfaces on Virtual Chassis member switches as VCPs and interconnect the members to form a LAG, be sure you have:

1. Installed the uplink modules in the SWA-0, SWA-1, SWA-2, and SWA-3 switches. See Installing an Uplink Module in a J-EX4200 Switch.
2. Powered on SWA-0, connected it to the network, and run the EZSetup program. See “Connecting and Configuring a J-EX Series Switch (CLI Procedure)” on page 161 or “Connecting and Configuring a J-EX Series Switch (J-Web Procedure)” on page 163 for details.
3. Configured SWA-0 with the virtual management Ethernet (VME) interface for remote, out-of-band management of the Virtual Chassis configuration, if desired. See “Configuring the Virtual Management Ethernet Interface for Global Management of a Virtual Chassis (CLI Procedure)” on page 797.
4. Ensured that SWA-1 is not powered on and then interconnected SWA-0 and SWA-1 using the dedicated VCPs on the rear panel.

   NOTE: The interfaces for the dedicated VCPs are operational by default. They do not need to be configured.

5. Ensured that SWA-2, SWA-3, and SWA-4 are not powered on. They are not connected in any way, so when initially powered up they will be standalone switches.

Overview and Topology

In this example, five J-EX4200 switches will be interconnected to form LAGs for ease of monitoring and manageability. Two of these switches (SWA-0 and SWA-1) are located in wiring closet A and the three others (SWA-2, SWA-3, and SWA-4) are located in wiring closet B. SWA-0 will form one LAG with SWA-2 and another LAG with SWA-4, and SWA-1 will form a LAG with SWA-3.

We will use fiber-optic cables connected to the uplink and network VCPs to interconnect the member switches in wiring closet A to the member switches in wiring closet B.

We will specify the highest mastership priority value (255) for SWA-0 to make it the master before we power on SWA-1. Because SWA-0 and SWA-1 are interconnected with the dedicated VCPs, the master detects that SWA-1 is a member of its Virtual Chassis configuration and assigns it a member ID.

We will use SWA-2 as the backup of the Virtual Chassis configuration. We will configure the same mastership priority value for SWA-2 (255) that we configured for the master. Because we power on SWA-0 before we power on SWA-2, SWA-0 retains mastership of the Virtual Chassis configuration.

NOTE: We recommend setting identical mastership priority values for the master and backup members for high availability and smooth transition of mastership in case the original master becomes unavailable.

We will configure the uplink module interfaces on three of the switches as uplink VCPs. On the J-EX4200-24F switch we will configure two of the network interfaces as VCPs. We will interconnect two of the uplink VCPs on SWA-0 with two of the uplink VCPs on SWA-2. Similarly, we will interconnect the two uplink VCPs on SWA-1 with the two uplink VCPs on SWA-3. Finally, we will connect the two remaining uplink VCPs on SWA-0 with two network VCPs on SWA-4. As a result, three LAGs will be automatically formed.

Figure 21 on page 771 shows the interconnections used to form LAGs using uplink VCPs and the network VCPs after the procedure below has been completed.

Figure 21: Virtual Chassis Interconnected Across Wiring Closets to Form LAGs

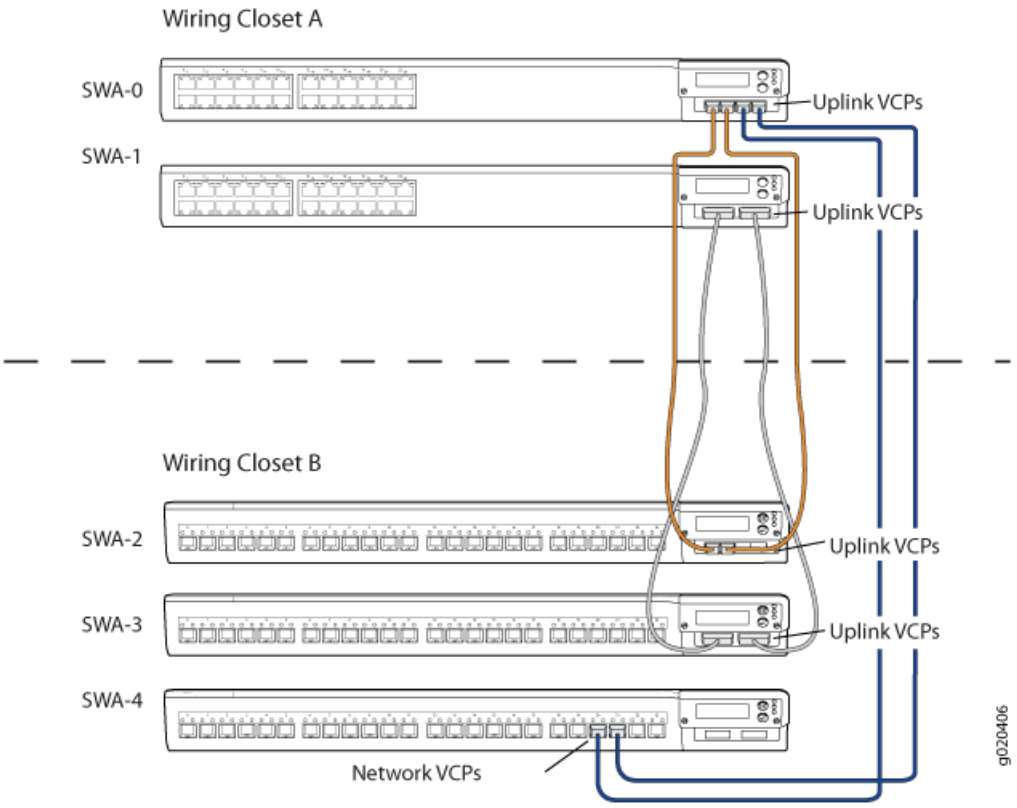

Configuration

To configure the Virtual Chassis uplink module interfaces and network interfaces as uplink VCPs and interconnect them between two wiring closets to form LAGs, perform these tasks:

Step-by-Step Procedure

To configure a Virtual Chassis across multiple wiring closets and interconnect them to form LAGs:

1. Configure the mastership priority of SWA-0 (member **0**) to be the highest possible value (**255**), thereby ensuring that it functions as the master of the expanded Virtual Chassis configuration:

```
[edit virtual-chassis]
user@SWA-0# set member 0 mastership-priority 255
```

2. Power on SWA-1.

3. Prepare the members in wiring closet A for interconnecting with the member switches in wiring closet B by setting all the uplink module interfaces on SWA-0 and two of the uplink module interfaces on SWA-1 as uplink VCPs:

```
user@SWA-0> request virtual-chassis vc-port set pic-slot 1 port 0
user@SWA-0> request virtual-chassis vc-port set pic-slot 1 port 1
user@SWA-0> request virtual-chassis vc-port set pic-slot 1 port 2
user@SWA-0> request virtual-chassis vc-port set pic-slot 1 port 3
user@SWA-0> request virtual-chassis vc-port set pic-slot 1 port 0 member 1
user@SWA-0> request virtual-chassis vc-port set pic-slot 1 port 1 member 1
```

NOTE: This example omits the specification of the member *member-id* option in configuring the uplink VCPs for SWA-0 (and, later, for SWA-2). The command applies by default to the switch where it is executed.

4. Power on SWA-2.

5. If SWA-2 was previously configured, revert to the factory default configuration.

6. Prepare SWA-2 in wiring closet B by configuring its mastership priority to be the highest possible value (255). Its member ID is currently **0**, because it is not yet interconnected with the other members of the Virtual Chassis configuration. It is operating as a standalone switch. Its member ID will change when it is interconnected.

```
[edit virtual-chassis]
user@SWA-2# set member 0 mastership-priority 255
```

NOTE: SWA-2 is configured with the same mastership priority value that we configured for SWA-0. However, the longer uptime of SWA-0 ensures that, once the interconnection is made, SWA-0 functions as the master and SWA-2 functions as the backup.

7. Specify two of the SFP uplink module interfaces in SWA-2 as uplink VCPs. The member IDs are **0**, because they are not yet interconnected with the other members of the Virtual Chassis configuration:

NOTE: The setting of the uplink VCPs remain intact when SWA-2 reboots and joins the Virtual Chassis configuration as member 2.

```
user@SWA-2> request virtual-chassis vc-port set pic-slot 1 port 0
user@SWA-2> request virtual-chassis vc-port set pic-slot 1 port 1
```

8. Power down SWA-2.
9. Physically interconnect SWA-0 and SWA-2 across wiring closets using two of the uplink VCPs on each switch.
10. Power on SWA-2. SWA-2 joins the Virtual Chassis configuration and a LAG is automatically formed between SWA-0 and SWA-2. In addition, although SWA-0 and SWA-2 have the same mastership priority value (255), SWA-0 was powered on first and thus has longer uptime. This results in SWA-0 retaining mastership while SWA-2 reboots and joins the now expanded Virtual Chassis configuration as the backup, with member ID **2**.
11. Power on SWA-3.
12. If SWA-3 was previously configured, revert to the factory default configuration.
13. Specify both uplink module interfaces in SWA-3 as uplink VCPs:

```
user@SWA-3> request virtual-chassis vc-port set pic-slot 1 port 0
user@SWA-3> request virtual-chassis vc-port set pic-slot 1 port 1
```

14. Power down SWA-3.
15. Physically interconnect SWA-3 with SWA-2 using their dedicated VCPs.
16. Physically interconnect SWA-1 and SWA-3 across wiring closets using their uplink VCPs.
17. Power on SWA-3. It joins the Virtual Chassis configuration as member **3**.

NOTE: Member ID 3 is assigned to SWA-3 because SWA-3 was powered on after members 0, 1, and 2.

A LAG is automatically formed between SWA-1 and SWA-3. In addition, both SWA-1 and SWA-3 have the default mastership priority value (128) and function in a linecard role.

18. Power on SWA-4.
19. If SWA-4 was previously configured, revert to the factory default configuration.
20. Configure two of the network interfaces on SWA-4 as uplink VCPs:

```
user@SWA-4> request virtual-chassis vc-port set pic-slot 0 port 20
user@SWA-4> request virtual-chassis vc-port set pic-slot 0 port 21
```

21. Power down SWA-4.

22. Physically interconnect SWA-4 and SWA-0 across wiring closets using the network VCPs on SWA-4 and the two remaining SFP uplink VCPs on SWA-0.
23. Power on SWA-4. A LAG is automatically formed between SWA-4 and SWA-0. In addition, SWA-4 joins the Virtual Chassis configuration in the linecard role.

Results Display the results of the configuration on SWA-0:

```
user@SWA-0> show configuration virtual-chassis
  member 0 {
    mastership-priority 255;
  }
  member 1 {
    mastership-priority 128;
  }
  member 2 {
    mastership-priority 255;
  }
  member 3 {
    mastership-priority 128;
  }
  member 4 {
    mastership-priority 128;
  }
}
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying the Member IDs and Roles of the Member Switches on page 774
- Verifying That the VCPs Are Operational on page 775

Verifying the Member IDs and Roles of the Member Switches

Purpose Verify that all the interconnected member switches are included within the Virtual Chassis configuration and that their roles are assigned appropriately.

Action Display the members of the Virtual Chassis configuration:

```
user@SWA-0> show virtual-chassis status

Virtual Chassis ID: 0000.e255.00e0

                                          Mastership            Neighbor List
Member ID  Status  Serial No   Model     Priority  Role        ID Interface

0 (FPC 0)  Prsnt   abc123      ex4200-48t    255   Master*     1 vcp-0
                                                               1 vcp-1
                                                               2 vcp-255/1/0
                                                               2 vcp-255/1/1
                                                               4 vcp-255/0/20
                                                               4 vcp-255/0/21

1 (FPC 1)  Prsnt   def456      ex4200-24t    128   Linecard    0 vcp-0
                                                               0 vcp-1
```

```
                                                                3 vcp-255/1/0
                                                                3 vcp-255/1/1

2 (FPC 2)   Prsnt   ghi789      ex4200-48t    255   Backup      3 vcp-0
                                                                3 vcp-1
                                                                0 vcp-255/1/0
                                                                0 vcp-255/1/1

3 (FPC 3)   Prsnt   jkl012      ex4200-24t    128   Linecard    2 vcp-0
                                                                2 vcp-1
                                                                1 vcp-255/1/0
                                                                1 vcp-255/1/1

4 (FPC 4)   Prsnt   mno345      ex4200-24f    128   Linecard    0 vcp-255/1/2
                                                                0 vcp-255/1/3
```

Meaning The **show virtual-chassis status** command lists the member switches interconnected in a Virtual Chassis configuration with the member IDs that have been assigned by the master, the mastership priority values, and the roles. It also displays the neighbor members with which each member is interconnected by the dedicated VCPs, by uplink VCPs, and by network VCPs.

Verifying That the VCPs Are Operational

Purpose Verify that the dedicated VCPs interconnecting member switches in wiring closets A and B and the uplink and network VCPs interconnecting the member switches between wiring closets are operational.

Action Display the Virtual Chassis interfaces:

```
user@SWA-0> show virtual-chassis vc-port all-members

fpc0:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status       Speed        Neighbor
or                             ID                (mbps)       ID  Interface
PIC / Port
vcp-0       Dedicated           1    Up           32000        1   vcp-0
vcp-1       Dedicated           2    Up           32000        1   vcp-1
1/0         Configured          3    Up           1000         2   vcp-255/1/0
1/1         Configured          3    Up           1000         2   vcp-255/1/1
1/2         Configured          4    Up           1000         4   vcp-255/0/20
1/3         Configured          4    Up           1000         4   vcp-255/0/21

fpc1:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status       Speed        Neighbor
or                             ID                (mbps)       ID  Interface
PIC / Port
vcp-0       Dedicated           1    Up           32000        0   vcp-0
vcp-1       Dedicated           2    Up           32000        0   vcp-1
1/0         Configured          3    Up           10000        3   vcp-255/1/0
1/1         Configured          3    Up           10000        3   vcp-255/1/1

fpc2:
--------------------------------------------------------------------------
```

```
Interface   Type              Trunk  Status       Speed        Neighbor
or                             ID                (mbps)       ID  Interface
PIC / Port
vcp-0       Dedicated           1    Up           32000        3   vcp-0
vcp-1       Dedicated           2    Up           32000        3   vcp-1
1/0         Configured          3    Up           1000         0   vcp-255/1/0
1/1         Configured          3    Up           1000         0   vcp-255/1/1
1/2                            -1    Down         1000
1/3                            -1    Down         1000

fpc3:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status       Speed        Neighbor
or                             ID                (mbps)       ID  Interface
PIC / Port
vcp-0       Dedicated           1    Up           32000        2   vcp-0
vcp-1       Dedicated           2    Up           32000        2   vcp-1
1/0         Configured          3    Up           10000        1   vcp-255/1/0
1/1         Configured          3    Up           10000        1   vcp-255/1/1

fpc4:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status       Speed        Neighbor
or                             ID                (mbps)       ID  Interface
PIC / Port
vcp-0       Dedicated           1    Down         32000
vcp-1       Dedicated           2    Down         32000
0/20        Configured          3    Up           1000         0   vcp-255/1/2
0/21        Configured          3    Up           1000         0   vcp-255/1/3
```

Meaning The dedicated VCPs are displayed as **vcp-0** and **vcp-1**. The uplink module interfaces that have been set as uplink VCPs are displayed as **1/0**, **1/1**, **1/2**, and **1/3**. The network interfaces that have been set as VCPs are displayed as **0/20** and **0/21**. The neighbor interface names of uplink and network VCPs are of the form **vcp-255/*pic/port***—for example, **vcp-255/1/0**. In that name, **vcp-255** indicates that the interface is a VCP, **1** is the uplink PIC number, and **0** is the port number. The **fpc** number is the same as the member ID. The trunk ID is a positive number ID assigned to the LAG formed by the Virtual Chassis. If no LAG is formed, the value is **–1**.

NOTE: Each switch assigns the trunk IDs to its local interfaces. As a result, the pair of interfaces that form one end of a LAG on one switch will have the same trunk ID, and the pair of interfaces that form the other end of the LAG will have the same trunk ID, but the trunk IDs on either end of the LAG might be different. For example, in Figure 21 on page 771, the uplink VCPs 1/2 and 1/3 on SWA-0 form a LAG with the network VCPs 0/20 and 0/21 on SWA-4. Uplink VCPs 1/2 and 1/3 on SWA-0 both have trunk ID 4, while network VCPs 0/20 and 0/21 on SWA-4 both have trunk ID 3. The trunk IDs are different between the switches because SWA-0 assigns the trunk IDs for its local uplink VCPs and SWA-4 assigns the trunk IDs for its local VCPs.

Troubleshooting

To troubleshoot a Virtual Chassis configuration that is interconnected across wiring closets, perform this task:

Troubleshooting Nonoperational VCPs

Problem An uplink VCP shows a status of **down**.

Solution

- Check the cable to make sure that it is properly and securely connected to the interfaces.
- If the VCP is an uplink module interface, make sure that it has been explicitly set as an uplink VCP.
- If the VCP is an uplink module interface, make sure that you have specified the options (***pic-slot***, ***port***, and ***member***) correctly.

Related Documentation

- Example: Configuring a Virtual Chassis with a Master and Backup in a Single Wiring Closet on page 717
- Example: Expanding a Virtual Chassis Configuration in a Single Wiring Closet on page 722
- Example: Setting Up a Multimember Virtual Chassis Access Switch with a Default Configuration on page 727
- Setting an Uplink Module Port as a Virtual Chassis Port (CLI Procedure) on page 792
- Reverting to the Default Factory Configuration for the J-EX Series Switch on page 341

Example: Configuring Automatic Software Update on Virtual Chassis Member Switches

The automatic software update feature automatically updates the Junos OS version on prospective member switches as they are added to a Virtual Chassis configuration of J-EX4200 Ethernet Switches so the new member switch immediately joins the Virtual Chassis configuration and is put in the active state. If the software version on the new switch is not the same as the version running on the master, the master keeps the new switch in the inactive state. If you have not enabled the automatic software update feature, you will have to manually install the correct software version on each prospective member switch as it is added to the Virtual Chassis configuration.

This example describes how to configure the Virtual Chassis automatic software update feature:

Requirements

This example uses the following hardware and software components:

- Three J-EX4200 switches

Before you begin, be sure you have:

1. Ensured that two member switches are running the same version of the Junos OS for J-EX Series switches so that they can form the initial Virtual Chassis configuration.
2. Cabled and powered on those two switches to create the Virtual Chassis configuration. See Connecting a Virtual Chassis Cable to a J-EX4200 Switch.
3. Ensured that you know the name or the URL of the software package to be used by the automatic software update feature.

Overview and Topology

For a standalone J-EX4200 switch to join an existing Virtual Chassis configuration, it must be running the same version of the Junos OS that is running on the Virtual Chassis master. If the software version on the new switch is not the same as the version running on the master, the master keeps the new switch in the inactive state.

The topology for this example consists of three J-EX Series switches. Two of the switches are connected in a Virtual Chassis configuration and are therefore running the same version of the Junos OS for J-EX Series switches. The third switch is a standalone switch that is running a different software version than the Virtual Chassis member switches. In this example, we will enable the automatic software update feature on the Virtual Chassis configuration and then add the third switch to the configuration. The master will detect the presence of the new switch, check the software version running on the new switch, and, because it is not the same version currently running on the master, will update the software version on the new switch and reboot the switch so that it can join the Virtual Chassis configuration and immediately be put in the active state.

Configuration

To configure automatic software update, perform this task:

Step-by-Step Procedure

To configure automatic software update:

1. Enable automatic software update and configure the path to the software package:

```
[edit]
user@switch# set virtual-chassis auto-sw-update package-name
/var/tmp/jinstall-ex-4200-10.2R1.1-domestic-signed.tgz
```

2. Connect and power on the new switch to be added to the existing Virtual Chassis configuration.

Results

Check the results of the configuration:

```
[edit virtual-chassis]
```

```
user@switch# show
auto-sw-update {
   package-name /var/tmp/jinstall-ex-4200-10.2R1.1-domestic-signed.tgz;
}
```

Verification

To verify that the software version on the new switch has been updated and that the switch has joined the Virtual Chassis configuration, perform this task:

- Verifying That the Software Version Is Updated on page 779

Verifying That the Software Version Is Updated

Purpose Verify that the new switch has joined the Virtual Chassis configuration.

NOTE: If the software version on the new switch had not been updated successfully, the master would not allow the switch to join the Virtual Chassis configuration.

Action

1. Issue the **show virtual-chassis status** command.
2. Check to see that the new member switch has been added.

```
user@switch> show virtual-chassis status
Virtual Chassis ID: 0019.e250.47a0
                                                Mastership            Neighbor List
Member ID  Status   Serial No    Model    priority    Role      ID  Interface
0 (FPC 0)  Prsnt    AK0207360276 ex4200-24t      255  Master*    1  vcp-1
                                                                 2  vcp-0
1 (FPC 1)  Prsnt    AK0207360281 ex4200-24t      255  Backup     2  vcp-1
                                                                 0  vcp-0
2 (FPC 2)  Prsnt    AJ0207391130 ex4200-48t      128  Linecard   0  vcp-1
                                                                 1  vcp-0
```

Meaning Because in the initial two-member Virtual Chassis configuration member 0 was the master and member 1 was the backup, the output shows that the new switch has been assigned member ID **2** and has been given the **Linecard** role. The **Status** field shows that member 2 is **Prsnt**, which means that it is in the active state.

Related Documentation

- Configuring Automatic Software Update on Virtual Chassis Member Switches (CLI Procedure) on page 800
- Adding a New Switch to an Existing Virtual Chassis Configuration (CLI Procedure) on page 786

CHAPTER 45

Configuring Virtual Chassis

Configuring a Virtual Chassis (CLI Procedure)

To take advantage of the scalability features of J-EX4200 switches, you can configure a Virtual Chassis that includes up to 10 member switches. You can interconnect the member switches using the dedicated Virtual Chassis ports (VCPs) on the back of the switch. You do not have to configure the interface for the dedicated VCPs. If you want to interconnect member switches that are located in different racks or wiring closets, interconnect them using uplinks configured as VCP interfaces. See “Setting an Uplink Module Port as a Virtual Chassis Port (CLI Procedure)” on page 792.

NOTE: A multimember Virtual Chassis configuration has two Routing Engines, one in the master and the other in the backup. Therefore, we recommend that you always use **commit synchronize** rather than simply **commit** to save configuration changes made for a Virtual Chassis. This ensures that the configuration changes are saved in both Routing Engines.

A Virtual Chassis can be configured with either:

- preprovisioned configuration—Allows you to deterministically control the member ID and role assigned to a member switch by tying it to its serial number.
- nonprovisioned configuration—The master sequentially assigns a member ID to other member switches. The role is determined by the mastership priority value and other factors in the master election algorithm.

- Configuring a Virtual Chassis with a Preprovisioned Configuration File on page 782
- Configuring a Virtual Chassis with a Nonprovisioned Configuration File on page 783

Configuring a Virtual Chassis with a Preprovisioned Configuration File

To configure a Virtual Chassis using a preprovisioned configuration:

1. Make a list of the serial numbers of all the switches to be connected in a Virtual Chassis configuration.
2. Note the desired role (**routing-engine** or **linecard**) of each switch. If you configure the member with a **routing-engine** role, it is eligible to function as a master or backup. If you configure the member with a **linecard** role, it is not eligible to become a master or backup.
3. Interconnect the member switches using the dedicated VCPs on the rear panel of switches. See Connecting a Virtual Chassis Cable to a J-EX4200 Switch.

   NOTE: Arrange the switches in sequence, either from top to bottom or from bottom to top (0–9).

4. Power on only the switch that you plan to use as the master switch (SWA-0). Do not power on the other switches at this time.
5. Run the EZ Setup program on SWA-0, specifying the identification parameters. See “Connecting and Configuring a J-EX Series Switch (CLI Procedure)” on page 161 for details.

   NOTE: The properties that you specify for SWA-0 apply to the entire Virtual Chassis configuration, including all the member listed in the preprovisioned configuration file.

6. Configure SWA-0 with the virtual management Ethernet (VME) interface for out-of-band management of the Virtual Chassis configuration, if desired.

```
[edit]
user@SWA-0# set interfaces vme unit 0 family inet address /ip-address/mask/
```

7. Specify the preprovisioned configuration mode:

```
[edit virtual-chassis]
user@SWA-0# set preprovisioned
```

8. Specify all the members that you want to included in the Virtual Chassis configuration, listing each switch's serial number with the desired member ID and the desired role:

```
[edit virtual-chassis]
user@SWA-0# set member 0 serial-number abc123 role routing-engine
user@SWA-0# set member 1 serial-number def456 role linecard
user@SWA-0# set member 2 serial-number ghi789 role linecard
user@SWA-0# set member 3 serial-number jkl012 role linecard
user@SWA-0# set member 4 serial-number mno345 role linecard
user@SWA-0# set member 5 serial-number pqr678 role routing-engine
user@SWA-0# set member 6 serial-number stu901 role linecard
user@SWA-0# set member 7 serial-number vwx234 role linecard
user@SWA-0# set member 8 serial-number yza567 role linecard
user@SWA-0# set member 9 serial-number bcd890 role linecard
```

9. Power on the member switches.

NOTE: You cannot modify the **mastership-priority** when you are using a preprovisioned configuration. The mastership priority values are generated automatically and controlled by the role that is assigned to the member switch in the configuration file. The two routing engines are assigned the same mastership priority value. However, the member that was powered on first has higher prioritization according to the master election algorithm. See "Understanding How the Master in a Virtual Chassis Configuration Is Elected" on page 698.

Configuring a Virtual Chassis with a Nonprovisioned Configuration File

To configure the Virtual Chassis using a nonprovisioned configuration:

1. Interconnect the member switches using the dedicated VCPs on the rear panel of switches. See Connecting a Virtual Chassis Cable to a J-EX4200 Switch.

NOTE: Arrange the switches in sequence, either from top to bottom or from bottom to top (0–9).

2. Power on only the switch that you plan to use as the master switch (SWA-0). Do not power on the other switches at this time.
3. Run the EZ Setup program on SWA-0, specifying the identification parameters. See "Connecting and Configuring a J-EX Series Switch (CLI Procedure)" on page 161 for details.

NOTE: The properties that you specify for SWA-0 apply to the entire Virtual Chassis configuration, including all the members interconnected through VCPs..

4. Configure SWA-0 with the virtual management Ethernet (VME) interface for out-of-band management of the Virtual Chassis configuration, if desired.

```
[edit]
user@SWA-0# set interfaces vme unit 0 family inet address /ip-address/mask/
```

5. Configure mastership priority for the master, backup, and other members, if desired:

```
[edit virtual-chassis]
user@SWA-0# set member 0 mastership-priority 255
user@SWA-0# set member 5 mastership-priority 255
```

6. Power on the member switches in sequential order, one by one.

NOTE: If you do not edit the Virtual Chassis configuration file, a nonprovisioned configuration is generated by default. The mastership priority value for each member switch is 128. The master role is selected by default. You can change the role that is performed by the members by modifying the **mastership-priority**. See "Configuring Mastership of the Virtual Chassis (CLI Procedure)" on page 790. We recommend that you specify the same mastership priority value for the desired master and backup members. We have assigned the highest possible mastership priority to two members. However, the member that was powered on first has higher prioritization according to the master election algorithm. See "Understanding How the Master in a Virtual Chassis Configuration Is Elected" on page 698. We have allowed the other members to use the default mastership priority, which qualifies them to function in the role of linecard.

NOTE: If you want to change the member ID that the master has assigned to a member switch, use the **request virtual-chassis renumber** command.

Related Documentation

- Configuring a Virtual Chassis (J-Web Procedure) on page 784
- Configuring Mastership of the Virtual Chassis (CLI Procedure) on page 790
- Setting an Uplink Module Port as a Virtual Chassis Port (CLI Procedure) on page 792
- Monitoring Virtual Chassis Configuration Status and Statistics on page 809

Configuring a Virtual Chassis (J-Web Procedure)

To take advantage of the scalability features of J-EX4200 switches, you can configure a Virtual Chassis that includes up to 10 member switches. You can interconnect the

member switches using the dedicated Virtual Chassis ports (VCPs) on the back of the switch. You do not have to configure the interface for the dedicated VCPs. If you want to interconnect member switches that are located in different racks or wiring closets, interconnect them using uplinks configured as VCP interfaces. See “Setting an Uplink Module Port as a Virtual Chassis Port (CLI Procedure)” on page 792.

To configure a Virtual Chassis for J-EX Series switches using the J-Web interface:

1. Select **Configure** > **Virtual Chassis**.

   

   NOTE: After you make changes to the configuration in this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See “Using the Commit Options to Commit Configuration Changes (J-Web Procedure)” on page 334 for details about all commit options.

2. The properties you can configure are displayed.

   The first section of the Virtual Chassis Configuration page displays the Virtual Chassis member configuration. the display includes a list of member switches, their member IDs, and the mastership priority.

   The second section displays the operational status of the Virtual Chassis configuration, member details, and the dedicated and configured Virtual Chassis ports (VCPs).

3. Enter information into the page as described in Table 107 on page 785.

4. Click one:

   - **Add**—To add a member's configuration to the Virtual Chassis configuration, click **Add**.
   - **Edit**—To modify an existing member's configuration, click **Edit**.
   - **Delete**—To delete the configuration of a member, click **Delete**.

5. To configure an uplink as a VCP, select the member in the Virtual Chassis members list and select **Action > Select Uplink Port as VCP**. Select the port from the list.

6. To delete an uplink VCP from a member, select the member in the Virtual Chassis members list and select **Action > Delete Uplink Port as VCP**.

Table 107: Virtual Chassis Configuration Fields

Field	Function	Your Action
Member Details		
Member ID	Specifies the identifier for the member switch. The master switch assigns member IDs.	Select an identifier from the list. Select an ID from **0** through **9**.
Priority	Specifies the mastership priority to be assigned to the member.	Select a number from **1** through **255**, with 255 being the highest priority (**128** is the default).

Table 107: Virtual Chassis Configuration Fields *(continued)*

Field	Function	Your Action
Disable Management VLAN	If you want to reserve an individual member's management Ethernet port for local troubleshooting, you can remove that port from being part of the Virtual Management Ethernet (VME).	Click to disable management VLAN on the port.
Refresh	Refreshes the operational status of Virtual Chassis members.	Click to refresh the operational status.

Related Documentation

- Configuring a Virtual Chassis (CLI Procedure) on page 781
- Example: Configuring a Virtual Chassis with a Master and Backup in a Single Wiring Closet on page 717
- Example: Configuring a Virtual Chassis Interconnected Across Multiple Wiring Closets on page 733
- Monitoring Virtual Chassis Configuration Status and Statistics on page 809
- Virtual Chassis Cabling Configuration Examples for J-EX4200 Switches
- Virtual Chassis Overview on page 691

Adding a New Switch to an Existing Virtual Chassis Configuration (CLI Procedure)

You can add one or more J-EX4200 switches to an existing Virtual Chassis configuration. Up to ten J-EX4200 switches can be included within a Virtual Chassis configuration. You can add the new switches to either type—nonprovisioned or preprovisioned—of Virtual Chassis configuration. See “Configuring a Virtual Chassis (CLI Procedure)” on page 781 for descriptions of these types.

To add a switch to an existing Virtual Chassis configuration, use the procedure that matches what you need to accomplish:

- Adding a New Switch to an Existing Virtual Chassis Configuration Within the Same Wiring Closet on page 786
- Adding a New Switch from a Different Wiring Closet to an Existing Virtual Chassis Configuration on page 787
- Adding a New Switch to an Existing Preprovisioned Virtual Chassis Configuration Using Autoprovisioning on page 789

Adding a New Switch to an Existing Virtual Chassis Configuration Within the Same Wiring Closet

Before you begin, be sure you have:

- Mounted the new switch in a rack.
- Confirmed that the new switch is powered off.

- If you are expanding a preprovisioned configuration, made a note of the serial number (on the back of the switch). You will need to edit the Virtual Chassis configuration to include the serial number of the new member switch.
- If you are expanding a preprovisioned configuration, edited the existing Virtual Chassis configuration to include the serial number of the new member switch. You can specify the role of the new member switch when you add its serial number in the Virtual Chassis configuration file. The parameters specified in the master Virtual Chassis configuration file are applied after the new member switch has been interconnected to an existing member switch.

NOTE: After you have created a preprovisioned Virtual Chassis configuration, you can use the autoprovisioning feature to add member switches to that configuration.

To add a new member switch to an existing Virtual Chassis configuration within the same wiring closet:

1. If the new member switch has been previously configured, revert that switch's configuration to the factory defaults. See "Reverting to the Default Factory Configuration for the J-EX Series Switch" on page 341.
2. Interconnect the unpowered new switch to at least one member of the existing Virtual Chassis configuration using the dedicated Virtual Chassis ports (VCPs).
3. Power on the new switch.
4. Confirm that the new member switch is now included within the Virtual Chassis configuration by checking the front-panel display for the member ID. It should display a member ID that is higher than 0 (1 through 9), because there is already at least one member of the Virtual Chassis configuration.

NOTE: If you are using a preprovisioned configuration, the member ID is automatically assigned to the member's serial number in the configuration file.

Adding a New Switch from a Different Wiring Closet to an Existing Virtual Chassis Configuration

To add a new switch from a different wiring closet to an existing Virtual Chassis configuration, you must use a long cable to connect the new member switch across wiring closets. You can use a port on an SFP or SFP+ uplink module, or an SFP network port on a J-EX4200-24F switch, and a fiber-optic cable for this purpose.

Before you begin, be sure you have:

- Installed the uplink modules needed for the Virtual Chassis configuration.
- Mounted the new switch in a rack.

- If the new member switch has been previously configured, reverted its configuration to the factory defaults. See “Reverting to the Default Factory Configuration for the J-EX Series Switch” on page 341.
- Powered on the new member switch as a standalone switch and configured its uplink module ports as VCPs. Otherwise, it cannot be recognized as a member switch by the master.
- If you are expanding a preprovisioned configuration, made a note of the serial number (on the back of the switch). You will need to edit the Virtual Chassis configuration to include the serial number of the new member switch.
- If you are expanding a preprovisioned configuration, edited the existing Virtual Chassis configuration to include the serial number of the new member switch. You can specify the role of the new member switch when you add its serial number in the Virtual Chassis configuration file. The parameters specified in the master Virtual Chassis configuration file are applied after the new member switch has been interconnected with its uplink VCP to an existing member switch.
- Confirmed that the new, currently standalone switch is powered off.
- Prepared an existing member switch for interconnecting with the new switch through an uplink module port by configuring an uplink module port as a VCP on the existing member switch.

NOTE: After you have created a preprovisioned Virtual Chassis configuration, you can use the autoprovisioning feature to add member switches to that configuration.

To add a new member switch that is going to be interconnected with the existing Virtual Chassis configuration across wiring closets:

1. Power on the new switch.
2. Connect a laptop or terminal to the console port of the switch, or use EZSetup on the standalone switch to specify temporary identification parameters. (When you interconnect the new member switch with the existing Virtual Chassis configuration, the master will overwrite and disable any specified parameters that conflict with the Virtual Chassis parameters or assigned member configuration.)
3. Use the CLI or the J-Web interface to set the uplink module ports as VCPs.

NOTE: If you are using a nonprovisioned configuration, you might configure the new member switch with a mastership priority value that is less than that of the existing member switches. Doing so ensures that the new member switch will function in a linecard role when it is included within the Virtual Chassis configuration.

4. Power off the new switch.

5. Interconnect the new member switch to at least one member of the existing Virtual Chassis configuration using the uplink module ports on each of the switches that have been configured as VCPs.
6. Power on the new member switch.
7. Confirm that the new member switch is now included within the Virtual Chassis configuration by checking the front-panel display for the member ID. It should display a member ID that is higher than 0 (1 through 9), because there is already at least one member of the Virtual Chassis configuration.

NOTE: If you are using a preprovisioned configuration, the member ID is automatically assigned to the member's serial number in the configuration file.

Adding a New Switch to an Existing Preprovisioned Virtual Chassis Configuration Using Autoprovisioning

Before you begin, be sure you have:

- Installed the uplink modules needed for the Virtual Chassis configuration.
- Mounted the new switch in a rack.
- Ensured that the preprovisioned Virtual Chassis configuration has an active master. For more information, see “Example: Configuring a Virtual Chassis Using a Preprovisioned Configuration File” on page 752.
- On the master, configured the Link Level Discovery Protocol (LLDP) on the uplink module ports that will be used as VCPs. LLDP is configured by default but might have been disabled. To configure LLDP, see “Configuring LLDP (CLI Procedure)” on page 2344 or “Configuring LLDP (J-Web Procedure)” on page 2345.
- Ensured that the new member switch has the factory-default configuration. If the new member switch has been previously configured, revert its configuration to the factory defaults. See “Reverting to the Default Factory Configuration for the J-EX Series Switch” on page 341.
- Made a note of the serial number (on the back of the switch). You will need to edit the Virtual Chassis configuration to include the serial number of the new member switch.
- Edited the existing Virtual Chassis preprovisioned configuration to include the serial number of the new member switch. You can specify the role of the new member switch when you add its serial number to the Virtual Chassis configuration file. The parameters specified in the master Virtual Chassis configuration file are applied to the new member switch after it has been interconnected through its uplink VCP to an existing member switch.
- Prepared an existing member switch to interconnect with the new switch through an uplink module port by configuring an uplink module port as a VCP on the existing member switch.

- Ensured that the operational modes of the uplink modules on the existing member switch and the new member switch match.
- Confirmed that the new member switch is powered off.
- Interconnected the existing switch with the new switch using the appropriate cable.

If these conditions are not met, autoprovisioning will not work and you will need to manually configure uplink module ports on the switch to be added to the configuration to be VCPs. For more information, see "Setting an Uplink Module Port as a Virtual Chassis Port (CLI Procedure)" on page 792.

To add a switch to an existing preprovisioned Virtual Chassis configuration using the autoprovisioning feature:

1. Power on the new member switch.
2. Confirm that the new member switch is now included in the Virtual Chassis configuration by checking the front-panel display for the member ID. It should display a member ID in the range from 0 through 9 because there was already at least one member of the Virtual Chassis configuration. The member ID is automatically assigned to the new member switch's serial number in the configuration file.

Related Documentation

- Example: Expanding a Virtual Chassis Configuration in a Single Wiring Closet on page 722
- Example: Setting Up a Multimember Virtual Chassis Access Switch with a Default Configuration on page 727
- Example: Configuring a Virtual Chassis Interconnected Across Multiple Wiring Closets on page 733
- Example: Configuring a Virtual Chassis Using a Preprovisioned Configuration File on page 752
- Example: Configuring Automatic Software Update on Virtual Chassis Member Switches on page 777
- Monitoring Virtual Chassis Configuration Status and Statistics on page 809
- Replacing a Member Switch of a Virtual Chassis Configuration (CLI Procedure) on page 811
- Reverting to the Default Factory Configuration for the J-EX Series Switch on page 341

Configuring Mastership of the Virtual Chassis (CLI Procedure)

You can designate the role (master, backup, or linecard) that a member switch performs within a Virtual Chassis configuration whether or not you are using a preprovisioned configuration.

NOTE: A multimember Virtual Chassis configuration has two Routing Engines, one in the master and the other in the backup. Therefore, we recommend that you always use **commit synchronize** rather than simply **commit** to save configuration changes made for a Virtual Chassis. This ensures that the configuration changes are saved in both Routing Engines.

This topic describes:

- Configuring Mastership Using a Preprovisioned Configuration File on page 791
- Configuring Mastership Using a Configuration File That Is Not Preprovisioned on page 792

Configuring Mastership Using a Preprovisioned Configuration File

To configure mastership using a preprovisioned configuration:

1. Note the serial numbers of the switches that you want to function in the master role and backup role.
2. Power on only the switch (SWA-0) that you want to function in the master role.
3. Edit the configuration to specify the preprovisioned configuration mode:

```
[edit virtual-chassis]
user@SWA-0# set preprovisioned
```

4. List the serial numbers of the member switches that you want to function as master and backup, specifying their role as **routing-engine**:

```
[edit]

user@SWA-0# set virtual-chassis member 0 serial-number abc123 role routing-engine

user@SWA-0# set virtual-chassis member 2 serial-number def456 role routing-engine
```

NOTE: You cannot directly modify the mastership priority value when you are using a preprovisioned configuration. The mastership priority values are generated automatically and controlled by the role that is assigned to the member switch in the configuration file. The two members assigned the routing-engine role are assigned the same mastership priority value (128). However, the member that was powered on first has higher prioritization according to the master election algorithm. See "Understanding How the Master in a Virtual Chassis Configuration Is Elected" on page 698. Only two members can be specified with the routing-engine role.

5. List the serial numbers of any other member switches that you want to include in the Virtual Chassis configuration. You may also specify their role as **linecard**, if desired.

Configuring Mastership Using a Configuration File That Is Not Preprovisioned

To configure mastership of the Virtual Chassis through a configuration that is not preprovisioned:

1. Power on only the switch that you want to function in the master role (SWA-0).
2. Configure the highest possible mastership priority value (**255**) for the member that you want to function in the master role:

```
[edit virtual-chassis]
user@SWA-0# set member 0 mastership-priority 255
```

3. Configure the same mastership priority value (continue to edit the Virtual Chassis configuration on the master) for the member that you want to be the backup (SWA-1):

```
[edit virtual-chassis]
user@SWA-0# set member 1 mastership-priority 255
```

NOTE: We recommend that the master and backup have the same mastership priority value to prevent the master and backup status from switching back and forth between master and backup members in failover conditions.

4. Use the default mastership priority value (**128**) for the remaining member switches or configure the mastership priority to a value that is lower than the value specified for members functioning in the master and backup roles.

Related Documentation

- Example: Configuring a Virtual Chassis Using a Preprovisioned Configuration File on page 752
- Example: Expanding a Virtual Chassis Configuration in a Single Wiring Closet on page 722
- Verifying the Member ID, Role, and Neighbor Member Connections of a Virtual Chassis Member on page 807
- Monitoring Virtual Chassis Configuration Status and Statistics on page 809
- Configuring a Virtual Chassis (CLI Procedure) on page 781
- Configuring a Virtual Chassis (J-Web Procedure) on page 784
- Understanding Virtual Chassis Configuration on page 704

Setting an Uplink Module Port as a Virtual Chassis Port (CLI Procedure)

You can interconnect J-EX4200 switches that are beyond the reach of the Virtual Chassis cables as members of a Virtual Chassis configuration by installing the optional SFP or SFP+ uplink module and connecting the uplink ports. You can also use the SFP network ports on a J-EX4200-24F for this purpose. To use the uplink ports or SFP network ports for interconnecting member switches, you must explicitly set the uplink ports as VCPs.

NOTE: When an uplink port is set as a VCP interface, it cannot be used for any other purpose. You can set one port as a VCP interface and configure the other port in trunk mode as an uplink to a distribution switch.

Before you set an uplink port as a VCP:

1. Install the uplink module in the member switches that you want to interconnect.
2. Power on and connect to the switch that you plan to designate as the master of the Virtual Chassis configuration.

NOTE: Do not power on the other switches at this point.

3. Run EZSetup on the switch that you are configuring to be the master. Follow the prompts to specify the hostname and other identification, time zone, and network properties. See “Connecting and Configuring a J-EX Series Switch (CLI Procedure)” on page 161 or “Connecting and Configuring a J-EX Series Switch (J-Web Procedure)” on page 163 for details. The properties that you specify for the master apply to the entire Virtual Chassis configuration, including all the member switches that you later interconnect with the master.
4. If you want to configure and manage the Virtual Chassis configuration remotely, specify the VME global management interface. You can configure the VME global management interface when you are setting up the master or you can do it after completing the other configuration steps for the Virtual Chassis. See “Configuring the Virtual Management Ethernet Interface for Global Management of a Virtual Chassis (CLI Procedure)” on page 797.
5. Configure mastership of the Virtual Chassis using either the nonprovisioned or preprovisioned configuration. See “Configuring Mastership of the Virtual Chassis (CLI Procedure)” on page 790 for details.

NOTE: A multimember Virtual Chassis configuration has two Routing Engines, one in the master and the other in the backup. Therefore, we recommend that you always use **commit synchronize** rather than simply **commit** to save configuration changes made for a Virtual Chassis configuration. This ensures that the configuration changes are saved in both Routing Engines.

To interconnect a Virtual Chassis configuration across longer distances, such as wiring closets, you need to:

- Prepare the existing Virtual Chassis configuration for interconnecting with a potential member switch that is beyond the reach of a Virtual Chassis cable by setting at least one uplink VCP on an existing member of Virtual Chassis configuration.
- Prepare the potential member switch for interconnecting with the existing Virtual Chassis configuration by setting at least one uplink VCP on the standalone switch.

NOTE: We recommend that you set two uplink VCPs within each wiring closet for redundancy.

This topic describes:

1. Setting an Uplink VCP Between Two Member Switches on page 794
2. Setting an Uplink VCP on a Standalone Switch on page 794

Setting an Uplink VCP Between Two Member Switches

Set an uplink port of a Virtual Chassis member as a VCP by executing the operational command **request virtual-chassis vc-port**.

NOTE: If you use the SFP+ uplink module, you must configure all member switches to support either 1-gigabit SFP transceivers or 10-gigabit SFP+ transceivers. See “Setting the Mode on an SFP+ Uplink Module (CLI Procedure)” on page 921.

To set the uplink ports for the local member switch (for example, member 0) and for a different member switch (for example, member 1) to function as VCPs:

1. Set one uplink port of member 0 as a VCP interface. You do not need to specify the **member** ***member-id*** option, because the command applies by default on the member where it is executed.

   ```
   user@SWA-0> request virtual-chassis vc-port set pic-slot 1 port 0
   ```

2. Set one uplink port of member 1 as a VCP interface.

   ```
   user@SWA-0>request virtual-chassis vc-port set pic-slot 1 port 0 member 1
   ```

   This example includes the member *member-id* option, because it is executed on a different member switch than the local member switch.

Setting an Uplink VCP on a Standalone Switch

To set an uplink VCP on a standalone switch, first power on the switch. You must set an uplink port on the standalone switch as a VCP prior to physically interconnecting the switch with the existing Virtual Chassis configuration. Otherwise, the master cannot detect that the switch is a member of the Virtual Chassis configuration.

To set one uplink VCP on the potential member (SWA-2), which is currently operating as a standalone switch:

1. Power on the standalone switch.
2. Set one uplink port as a VCP interface. You do not need to specify the **member** ***member-id*** option, because the command applies by default on the member where it is executed.

   ```
   user@SWA-2> request virtual-chassis vc-port set pic-slot 1 port 0
   ```

NOTE: If you do specify the **member** *member-id* option, use member ID 0. Because the switch is not yet interconnected with the other members of the Virtual Chassis configuration, its current member ID is 0. Its member ID will change when it is interconnected with the Virtual Chassis configuration. It does not impact the functioning of the uplink VCP that its VCP interface is set with 0 as the member ID. The VCP interface has significance only on the local switch.

3. After you have set the uplink VCP on the standalone switch, physically interconnect its uplink port with the VCP uplink ports of the members in the existing Virtual Chassis configuration.
4. The new member switch reboots and joins the now expanded Virtual Chassis configuration with a different member ID.

NOTE: The setting for the new member switch's uplink VCP remains intact and is not affected by the change of member ID.

5. If you have additional members in the second wiring closet, set a redundant VCP uplink on another member switch by issuing the **request virtual-chassis vc-port** command.

Related Documentation

- Configuring a Virtual Chassis (CLI Procedure) on page 781
- Configuring a Virtual Chassis (J-Web Procedure) on page 784
- Example: Configuring a Virtual Chassis Interconnected Across Multiple Wiring Closets on page 733
- Example: Configuring a Virtual Chassis Using a Preprovisioned Configuration File on page 752
- Monitoring Virtual Chassis Configuration Status and Statistics on page 809

Setting an Uplink Module Port or a J-EX4200-24F Network Port as a Virtual Chassis Port Using the LCD Panel

You can interconnect J-EX4200 switches that are beyond the reach of the Virtual Chassis cables as members of a Virtual Chassis configuration by installing the optional SFP or SFP+ uplink module and connecting the uplink module ports. You can also use the network ports on J-EX4200-24F switches to interconnect Virtual Chassis member switches. To use the uplink module ports or the J-EX4200-24F network ports for interconnecting member switches, you must explicitly set the ports as VCPs.

This topic describes how to set the uplink module ports and the J-EX4200-24F network ports as VCPs using the LCD panel on the front of J-EX4200 switches.

In this procedure, we show how to configure uplink module port **ge-0/1/2** as a VCP.

To set an uplink module port as a VCP using the LCD panel:

1. Press **Menu** until you see **MAINTENANCE MENU**.
2. Press **Menu** until you see **REQUEST VC PORT**.
3. Press **Enter**. You will see **SET VC PORT?**.
4. Press **Enter**. You will see **SET FPC 0?**.
5. Press **Enter**. You will see **SET PIC 0?**.
6. Press **Menu** until you see **SET PIC 1?**.
7. Press **Enter**. You will see **SET PORT 0?**.
8. Press **Menu** until you see **SET PORT 2?**.
9. Press **Enter**. You will see **CONFIGURING ....**
10. Once the configuration has been accepted, press **Enter** to return to the **MAINTENANCE** menu.

You can also use the LCD panel to delete a VCP, thus resetting the port to an uplink module port or a J-EX4200-24F network port.

To reset **vcp-0/1/2** to an uplink module port using the LCD panel:

1. Press **Menu** until you see **MAINTENANCE MENU**.
2. Press **Menu** until you see **REQUEST VC PORT**.
3. Press **Enter**. You will see **SET VC PORT?**.
4. Press **Menu**. You will see **DELETE VC PORT?**.
5. Press **Enter**. You will see **DELETE FPC 0?**.
6. Press **Enter**. You will see **DELETE PIC 0?**.
7. Press **Menu** until you see **DELETE PIC 1?**.
8. Press **Enter**. You will see **DELETE PORT 0?**.
9. Press **Menu** until you see **DELETE PORT 2?**.
10. Press **Enter**. You will see **CONFIGURING ....**
11. Once the configuration has been accepted, press **Enter** to return to the **MAINTENANCE** menu.

Related Documentation

- LCD Panel in J-EX4200 Switches
- Configuring a Virtual Chassis (CLI Procedure) on page 781
- Configuring a Virtual Chassis (J-Web Procedure) on page 784
- Setting an Uplink Module Port as a Virtual Chassis Port (CLI Procedure) on page 792
- Understanding Interface Naming Conventions on J-EX Series Switches on page 865

Configuring the Virtual Management Ethernet Interface for Global Management of a Virtual Chassis (CLI Procedure)

If you want to configure and manage a Virtual Chassis remotely through SSH or Telnet, configure the virtual management Ethernet (VME) interface on the master of the Virtual Chassis. You can configure and manage all members of the Virtual Chassis through this single global interface.

1. Power on the switch that you want to function as the master.
2. Check the front-panel LCD to confirm that the switch has powered on correctly.
3. Run the EZ Setup program on the switch, specifying the identification parameters. See "Connecting and Configuring a J-EX Series Switch (CLI Procedure)" on page 161 or "Connecting and Configuring a J-EX Series Switch (J-Web Procedure)" on page 163 for details.

To configure the VME:

```
[edit]
user@SWA-0# set interfaces vme unit 0 family inet address /ip-address/mask/
```

Related Documentation

- Example: Configuring a Virtual Chassis with a Master and Backup in a Single Wiring Closet on page 717
- Understanding Global Management of a Virtual Chassis Configuration on page 699

Configuring the Timer for the Backup Member to Start Using Its Own MAC Address, as Master of Virtual Chassis (CLI Procedure)

When a backup member takes control of a Virtual Chassis configuration because of a reset or other temporary failure, the backup uses the MAC address of the old master. This helps to ensure a smooth transition of mastership with no disruption to network connectivity.

The MAC persistence timer is used in situations when the master is no longer a member of the Virtual Chassis configuration, because it has been physically disconnected or removed. If the old master does not rejoin the Virtual Chassis configuration before the timer elapses, the new master starts using its own MAC address.

The default timer value is 10 minutes. There are no minimum or maximum limits.

Before you begin configuring the timer, ensure that you have at least two member switches in the Virtual Chassis configuration. To configure or modify the MAC persistence timer, use the following command:

```
[edit virtual-chassis]
user@switch# set mac-persistence-timer 30
```

This command modifies the MAC persistence timer value to specify a timer value of 30 minutes rather than the default timer value of 10 minutes.

Related Documentation

- Configuring a Virtual Chassis (CLI Procedure) on page 781
- Configuring a Virtual Chassis (J-Web Procedure) on page 784
- Understanding Virtual Chassis Components on page 694

Configuring Fast Failover in a Virtual Chassis Configuration

The Virtual Chassis fast failover feature is a hardware-assisted failover mechanism that automatically reroutes traffic and reduces traffic loss in the event of a link or switch failure. If a link between two members fails, traffic flow between those members must be rerouted quickly so that there is minimal traffic loss.

While fast failover is enabled by default on dedicated Virtual Chassis ports (VCPs), you must manually enable fast failover on uplink module ports that have been configured as VCPs.

Before you begin configuring fast failover, ensure that the dedicated VCPs or uplink module VCPs are connected in a ring topology.

- To reenable the fast failover feature on all dedicated VCPs in a ring:

```
[edit]
user@switch# delete virtual-chassis fast-failover vcp disable
```

- To configure the fast failover feature on all SFP uplink module VCPs in a ring:

```
[edit]
user@switch# set virtual-chassis fast-failover ge
```

Related Documentation

- Example: Configuring Fast Failover on Uplink Module VCPs to Reroute Traffic When a Virtual Chassis Member Switch or Inter-Member Link Fails on page 763
- Disabling Fast Failover in a Virtual Chassis Configuration on page 799
- Setting an Uplink Module Port as a Virtual Chassis Port (CLI Procedure) on page 792
- Configuring a Virtual Chassis (CLI Procedure) on page 781
- Configuring a Virtual Chassis (J-Web Procedure) on page 784
- Understanding Fast Failover in a Virtual Chassis Configuration on page 706

Disabling Fast Failover in a Virtual Chassis Configuration

While fast failover is enabled by default on dedicated Virtual Chassis ports (VCPs), you can manually disable fast failover on dedicated VCPs using the **set virtual-chassis fast-failover vcp disable** command.

- To disable the fast failover feature on all dedicated VCPs in a ring:

```
[edit]
user@switch# set virtual-chassis fast-failover vcp disable
```

- To disable the fast failover feature on all SFP uplink module VCPs in a ring:

```
[edit]
user@switch# delete virtual-chassis fast-failover ge
```

Related Documentation

- Example: Configuring Fast Failover on Uplink Module VCPs to Reroute Traffic When a Virtual Chassis Member Switch or Inter-Member Link Fails on page 763
- Configuring Fast Failover in a Virtual Chassis Configuration on page 798
- Setting an Uplink Module Port as a Virtual Chassis Port (CLI Procedure) on page 792
- Configuring a Virtual Chassis (CLI Procedure) on page 781
- Configuring a Virtual Chassis (J-Web Procedure) on page 784
- Understanding Fast Failover in a Virtual Chassis Configuration on page 706

Disabling Split and Merge in a Virtual Chassis Configuration (CLI Procedure)

The split and merge feature is enabled by default on J-EX4200 switches in a Virtual Chassis configuration. You can disable the split and merge feature using the **set virtual-chassis no-split-detection** command. If you disable the split and merge feature and the Virtual Chassis configuration splits, both parts of the split Virtual Chassis configuration remain active.

In a preprovisioned Virtual Chassis configuration, if both of the Routing Engines end up in the same Virtual Chassis configuration after a split, the other split Virtual Chassis configuration remains inactive. If the Routing Engines end up in different parts of the split Virtual Chassis configuration and the rest of the member switches are configured as having linecard roles, then a backup Routing Engine might not be selected for either part.

To disable the split and merge feature in a Virtual Chassis configuration:

```
[edit]
user@switch# set virtual-chassis no-split-detection
```

Related Documentation

- Example: Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge on page 767
- Configuring a Virtual Chassis (CLI Procedure) on page 781
- Configuring a Virtual Chassis (J-Web Procedure) on page 784

- Understanding Split and Merge in a Virtual Chassis Configuration on page 712
- Understanding Virtual Chassis Configuration on page 704

Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge (CLI Procedure)

Every Virtual Chassis configuration has a unique ID that is automatically assigned when the Virtual Chassis configuration is formed. You can also explicitly assign a Virtual Chassis ID using the **set virtual-chassis id** command. When two Virtual Chassis configurations attempt to merge, the Virtual Chassis ID that you assigned takes precedence over the automatically assigned Virtual Chassis IDs and becomes the ID for the newly merged Virtual Chassis configuration.

To configure the Virtual Chassis ID:

```
[edit]
user@switch# set virtual-chassis id id
```

Related Documentation

- Example: Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge on page 767
- Configuring a Virtual Chassis (CLI Procedure) on page 781
- Configuring a Virtual Chassis (J-Web Procedure) on page 784
- Understanding Split and Merge in a Virtual Chassis Configuration on page 712
- Understanding Virtual Chassis Configuration on page 704

Configuring Automatic Software Update on Virtual Chassis Member Switches (CLI Procedure)

The automatic software update feature allows you to automatically update the software version on prospective member switches as they are added so that they can join a Virtual Chassis configuration.

Before you begin, ensure that you know the name or the URL of the software package to be used by the automatic software update feature.

To configure the automatic software update feature:

```
[edit]
user@switch# set virtual-chassis auto-sw-update package-name package-name
```

If the software package is located on a local directory on the switch, use the following format for ***package-name***:

/pathname/package-name

If the software package is to be downloaded and installed from a remote location, use one of the following formats:

ftp://*hostname/pathname/package-name*

ftp://*username*:*prompt*@ftp.*hostname.net/package-name*

http://*hostname/pathname/package-name*

Related Documentation

- Example: Configuring Automatic Software Update on Virtual Chassis Member Switches on page 777
- Understanding Automatic Software Update on Virtual Chassis Member Switches on page 715

Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure)

In a Virtual Chassis configuration, one member switch is assigned the master role and has the master Rouing Engine. Another member switch is assigned the backup role and has the backup Routing Engine. Graceful Routing Engine switchover (GRES) enables the master and backup Routing Engines in a Virtual Chassis configuration to switch from the master to backup without interruption to packet forwarding. When you configure graceful Routing Engine switchover, the backup Routing Engine automatically synchronizes with the master Routing Engine to preserve kernel state information and the forwarding state.

To set up a Virtual Chassis configuration to use graceful Routing Engine switchover (GRES):

1. Set up a minimum of two J-EX4200 switches in a Virtual Chassis configuration with mastership priority of 255:

```
[edit]
user@switch# set virtual-chassis member 0 mastership-priority 255

[edit]
user@switch# set virtual-chassis member 1 mastership-priority 255
```

2. Set up graceful Routing Engine switchover:

```
[edit]
user@switch# set chassis redundancy graceful-switchover
```

Commit the configuration.

Related Documentation

- Example: Configuring a Virtual Chassis with a Master and Backup in a Single Wiring Closet on page 717
- Configuring a Virtual Chassis (CLI Procedure) on page 781
- Configuring a Virtual Chassis (J-Web Procedure) on page 784
- High Availability Features for J-EX Series Switches Overview on page 18
- Understanding Virtual Chassis Configuration on page 704
- For more information about graceful Routing Engine switchover, see the *Junos OS High Availability Configuration Guide* at http://www.juniper.net/techpubs/software/junos/index.html.

CHAPTER 46

Verifying Virtual Chassis Configuration

Command Forwarding Usage with a Virtual Chassis Configuration

Some CLI commands can be run either on all members or on a specific member of a Virtual Chassis configuration. This functionality is referred to as command forwarding.

For example, to collect information about your system prior to contacting Dell Support (see "Requesting Technical Support" on page lxxi), use the command **request support information all-members** to gather data for all the member switches. If you want to gather this data only for a particular member switch, use the command **request support information member** ***member-id*** .

Table 108 on page 804 provides a list of commands that can be run either on all members of the Virtual Chassis configuration or on a specific member switch.

Table 108: Commands That Can be Run on All or Specific Members of the Virtual Chassis Configuration

Commands Available for Command Forwarding	Purpose	all-members	member-*member-id*
request support information	Use this command when you contact Dell Support about your component problem. This command is the equivalent of using the following CLI commands: • **show version** • **show chassis firmware** • **show chassis hardware** • **show chassis environment** • **show interfaces extensive** (for each configured interface) • **show configuration** (excluding any SECRET-DATA) • **show system virtual-memory**	Displays information for all members of the Virtual Chassis configuration.	Displays information for the specified member switch.
request system partition hard-disk	Set up the hard disk for partitioning. After this command is issued, the hard disk is partitioned the next time the system is rebooted. When the hard disk is partitioned, the contents of **/altroot** and **/altconfig** are saved and restored. All other data on the hard disk is at risk of being lost.	Partitions the hard disk on all members of the Virtual Chassis configuration.	Partitions the hard disk on the specified member switch.
request system reboot	Reboot the Junos OS for J-EX Series switches after a software upgrade and occasionally to recover from an error condition.	Reboots all members of the Virtual Chassis configuration.	Reboots the specified member switch.
request system snapshot	Back up the currently running and active file system.	Backs up the file systems on all members of the Virtual Chassis configuration.	Backs up the file system on the specified member switch.
request system storage cleanup	Free storage space on the switch by rotating log files and proposing a list of files for deletion. User input is required for file deletion.	Runs cleanup on all members of the Virtual Chassis configuration.	Runs cleanup on the specified member switch.
show log user	Display users who are viewing the system log.	Displays information for all members of the Virtual Chassis configuration.	Displays information for the specified member switch.

Table 108: Commands That Can be Run on All or Specific Members of the Virtual Chassis Configuration *(continued)*

Commands Available for Command Forwarding	Purpose	all-members	member-*member-id*
show system alarms	Display active system alarms.	Displays information for all members of the Virtual Chassis configuration.	Displays information for the specified member switch.
show system audit	Display the state and checksum values for file systems.	Displays information for all members of the Virtual Chassis configuration.	Displays information for the specified member switch.
show system boot-messages	Display initial messages generated by the system kernel upon startup. These messages are the contents of **/var/run/dmesg.boot**.	Displays information for all members of the Virtual Chassis configuration.	Displays information for the specified member switch.
show system core-dumps	Display a core file generated by an internal Junos OS process.	Displays information for all members of the Virtual Chassis configuration.	Displays information for the specified member switch.
show system directory-usage	Display directory usage information.	Displays information for all members of the Virtual Chassis configuration.	Displays information for the specified member switch.
show system reboot	Display pending system reboots or halts.	Displays information for all members of the Virtual Chassis configuration.	Displays information for the specified member switch.
show system snapshot	Display information about the backup software that is located in the **/altroot** and **/altconfig** file systems. To back up software, use the request system snapshot command.	Displays information for all members of the Virtual Chassis configuration.	Displays information for the specified member switch.
show system software	Display the Junos OS extensions loaded on your switch.	Displays information for all members of the Virtual Chassis configuration.	Displays information for the specified member switch.
show system statistics	Display systemwide protocol-related statistics.	Displays information for all members of the Virtual Chassis configuration.	Displays information for the specified member switch.
show system storage	Display statistics about the amount of free disk space in the switch's file systems.	Displays information for all members of the Virtual Chassis configuration.	Displays information for the specified member switch.

Table 108: Commands That Can be Run on All or Specific Members of the Virtual Chassis Configuration *(continued)*

Commands Available for Command Forwarding	Purpose	all-members	member-*member-id*
show system uptime	Display the current time and information about how long the switch, the switch software, and any existing protocols have been running	Displays information for all members of the Virtual Chassis configuration.	Displays information for the specified member switch.
show system users	Show all users who are currently logged in.	Shows all users who are currently logged in to any members of the Virtual Chassis configuration.	Shows all users who are currently logged in to the specified member switch.
show system virtual-memory	Display the usage of the Junos OS kernel memory, listed first by size of allocation and then by type of usage. Use **show system virtual-memory** for troubleshooting with Dell Support (see "Requesting Technical Support" on page lxxi).	Displays information for all members of the Virtual Chassis configuration.	Displays information for the specified member switch.

Table 109 on page 806 shows a list of commands that are relevant only to the master. Do not use the options **all-members** or **member-*member-id*** with these commands.

Table 109: Commands Relevant Only to the Master

Commands Relevant Only to the Master	Purpose
set date	Set the data and time.
show system buffers	Display information about the buffer pool that the Routing Engine uses for local traffic. Local traffic is the routing and management traffic that is exchanged between the Routing Engine and the Packet Forwarding Engine within the switch, as well as the routing and management traffic from IP (that is, from OSPF, BGP, SNMP, ping operations, and so on).
show system connections	Display information about the active IP sockets on the Routing Engine. Use this command to verify which servers are active on a system and which connections are currently in progress.
show system processes	Display information about software processes that are running on the switch and that have controlling terminals.

Related Documentation

- Monitoring Virtual Chassis Configuration Status and Statistics on page 809
- Understanding Virtual Chassis Components on page 694

- *Junos OS System Basics and Services Command Reference* at http://www.juniper.net/techpubs/software/junos/

Verifying the Member ID, Role, and Neighbor Member Connections of a Virtual Chassis Member

Purpose You can designate the role that a member performs within a Virtual Chassis configuration or you can allow the role to be assigned by default. You can designate the member ID that is assigned to a specific switch by creating a permanent association between the switch's serial number and a member ID, using a preprovisioned configuration. Or you can let the member ID be assigned by the master, based on the sequence in which the member switch is powered on and on which member IDs are currently available.

The role and member ID of the member switch are displayed on the front-panel LCD.

Each member switch can be cabled to one or two other member switches, using either the dedicated Virtual Chassis ports (VCPs) on the rear panel, an uplink module port that has been configured as a VCP, or an SFP network port on a J-EX4200-24F switch that has been configured as a VCP. The members that are cabled together are considered neighbor members.

Action To display the role and member ID assignments using the CLI, use the **show virtual-chassis status** command:

```
user@SWA-0> show virtual-chassis status

Virtual Chassis ID: 0000.e255.00e0

                                          Mastership            Neighbor List
Member ID  Status  Serial No   Model     Priority  Role        ID, Interface

0 (FPC 0)  Prsnt   abc123      ex4200-48t    255  Master*     1 vcp-0
                                                              2 vcp-1

1 (FPC 1)  Prsnt   def456      ex4200-24t    255  Backup      2 vcp-0
                                                              0 vcp-1

2 (FPC 2)  Prsnt   abd231      ex4200-24t    128  Linecard    0 vcp-0
                                                              1 vcp-1
```

Meaning This output verifies that three J-EX4200 switches have been interconnected as a Virtual Chassis configuration using their dedicated VCPs . The display shows which of the VCPs is connected to which neighbor. The first port (**vcp-0**) of member **0** is connected to member **1** and the second port of member **0** (**vcp-1**) is connected to member **2**. The FPC slots for J-EX Series switches are the same as the member IDs.

The **Mastership Priority** values indicate that the master and backup members have been explicitly configured, because they are not using the default value (**128**).

Related Documentation

- Configuring Mastership of the Virtual Chassis (CLI Procedure) on page 790
- Configuring a Virtual Chassis (CLI Procedure) on page 781

- Configuring a Virtual Chassis (J-Web Procedure) on page 784
- Example: Expanding a Virtual Chassis Configuration in a Single Wiring Closet on page 722
- Example: Setting Up a Multimember Virtual Chassis Access Switch with a Default Configuration on page 727
- Monitoring Virtual Chassis Configuration Status and Statistics on page 809

Verifying That the Virtual Chassis Ports Are Operational

Purpose Use the **show virtual-chassis vc-port** command to display the status of Virtual Chassis ports (VCPs).

NOTE: The interfaces for VCPs are not displayed when you issue the **show interfaces ge-** command.

Action Display the VCPs:

```
user@SWA-0> show virtual-chassis vc-port all-members

fpc0:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status       Speed        Neighbor
or                             ID                (mbps)       ID  Interface
PIC / Port
vcp-0       Dedicated           1    Up           32000        1   vcp-0
vcp-1       Dedicated           2    Up           32000        1   vcp-1
1/0         Configured          3    Up           1000         2   vcp-255/1/0
1/1         Configured          3    Up           1000         2   vcp-255/1/1
1/2         Configured          4    Up           1000         4   vcp-255/0/20
1/3         Configured          4    Up           1000         4   vcp-255/0/21

fpc1:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status       Speed        Neighbor
or                             ID                (mbps)       ID  Interface
PIC / Port
vcp-0       Dedicated           1    Up           32000        0   vcp-0
vcp-1       Dedicated           2    Up           32000        0   vcp-1
1/0         Configured          3    Up           10000        3   vcp-255/1/0
1/1         Configured          3    Up           10000        3   vcp-255/1/1

fpc2:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status       Speed        Neighbor
or                             ID                (mbps)       ID  Interface
PIC / Port
vcp-0       Dedicated           1    Up           32000        3   vcp-0
vcp-1       Dedicated           2    Up           32000        3   vcp-1
1/0         Configured          3    Up           1000         0   vcp-255/1/0
1/1         Configured          3    Up           1000         0   vcp-255/1/1
1/2                            -1    Down         1000
1/3                            -1    Down         1000
```

```
fpc3:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status       Speed        Neighbor
or                             ID                (mbps)       ID  Interface
PIC / Port
vcp-0       Dedicated           1    Up           32000        2   vcp-0
vcp-1       Dedicated           2    Up           32000        2   vcp-1
1/0         Configured          3    Up           10000        1   vcp-255/1/0
1/1         Configured          3    Up           10000        1   vcp-255/1/1

fpc4:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status       Speed        Neighbor
or                             ID                (mbps)       ID  Interface
PIC / Port
vcp-0       Dedicated           1    Down         32000
vcp-1       Dedicated           2    Down         32000
0/20        Configured          3    Up           1000         0   vcp-255/1/2
0/21        Configured          3    Up           1000         0   vcp-255/1/3
```

Meaning The dedicated VCPs are displayed as **vcp-0** and **vcp-1**. The uplink module interfaces that have been set as uplink VCPs are displayed as **1/0**, **1/1**, **1/2**, and **1/3**. The J-EX4200-24F network interfaces that have been set as VCPs are displayed as **0/20** and **0/21**. The neighbor interface names of uplink and network VCPs are of the form **vcp-255/*pic/port***—for example, **vcp-255/1/0**. In that name, **vcp-255** indicates that the interface is a VCP, **1** is the uplink PIC number, and **0** is the port number. The **fpc** number is the same as the member ID. The trunk ID is a positive number ID assigned to the LAG formed by the Virtual Chassis. If no LAG is formed, the value is **–1**.

Related Documentation

- Monitoring Virtual Chassis Configuration Status and Statistics on page 809
- Configuring a Virtual Chassis (CLI Procedure) on page 781
- Configuring a Virtual Chassis (J-Web Procedure) on page 784
- Example: Configuring a Virtual Chassis Interconnected Across Multiple Wiring Closets on page 733

Monitoring Virtual Chassis Configuration Status and Statistics

Purpose Use the monitoring functionality to view the following information about Virtual Chassis members and ports:

- Member details and how members are connected with each other.
- Traffic statistics for Virtual Chassis ports of the selected members.
- Details of the Virtual Chassis port packet counters.

Action To view Virtual Chassis monitoring details in the J-Web interface, select **Monitor > Virtual Chassis**.

To view member details for all members in the CLI, enter the following command:

show virtual-chassis status

To view Virtual Chassis port traffic statistics for a specific member in the CLI, enter the following command:

show virtual-chassis vc-port statistics member ***member-id***

To view the path a packet takes when going from a source interface to a destination interface in a Virtual Chassis configuration using the CLI, enter the following command:

show virtual-chassis vc-path

Meaning In the J-Web interface the top half of the screen displays details of the Virtual Chassis configuration, such as:

- Member
- Role
- Interface
- Type
- Speed
- Neighboring Member ID
- Link Status
- Error count

Click the **Stop** button to stop fetching values from the switch, and click the **Start** button to start plotting data again from the point where it was stopped.

To view a graph of the statistics for the selected Virtual Chassis port of the member, click **Show Graph**.

Refresh Interval (sec)—Displays the time interval you have set for page refresh.

Click **Clear Statistics** to clear the monitoring statistics for the selected member switch. You can specify the interval at which the member details and statistics must be refreshed.

The bottom half of the screen displays a chart of the Virtual Chassis statistics, and the port packet counters.

For details about the output from CLI commands, see **show virtual-chassis status** and **show virtual-chassis vc-port statistics**.

Related Documentation

- Configuring a Virtual Chassis (CLI Procedure) on page 781
- Configuring a Virtual Chassis (J-Web Procedure) on page 784
- Example: Configuring a Virtual Chassis with a Master and Backup in a Single Wiring Closet on page 717

- Verifying the Member ID, Role, and Neighbor Member Connections of a Virtual Chassis Member on page 807

Replacing a Member Switch of a Virtual Chassis Configuration (CLI Procedure)

You can replace a member switch of a Virtual Chassis configuration without disrupting network service for the other members. You can retain the existing configuration of the member switch and apply it to a new member switch, or you can free up the member ID and make it available for assignment to a new member switch.

To replace a member switch, use the procedure that matches what you need to accomplish:

- Remove, Repair, and Reinstall the Same Switch on page 811
- Remove a Member Switch, Replace with a Different Switch, and Reapply the Old Configuration on page 812
- Remove a Member Switch and Make Its Member ID Available for Reassignment to a Different Switch on page 812

Remove, Repair, and Reinstall the Same Switch

If you need to repair a member switch, you can remove it from the Virtual Chassis configuration without disrupting network service for the other members. The master stores the configuration of the member ID so that it can be reapplied when the member switch (with the same base MAC address) is reconnected.

1. Power off and disconnect the member switch to be repaired.
2. Repair, as necessary.
3. Reconnect and power on the member switch.

Remove a Member Switch, Replace with a Different Switch, and Reapply the Old Configuration

If you are unable to repair a member switch, you can replace it with a different member switch and retain the old configuration. The master stores the configuration of the member that was removed. When you connect a different member switch, the master assigns a new member ID. But the old configuration is still stored under the previous member ID of the previous member switch.

NOTE: If you have used a preprovisioned configuration, use the replace command to change the serial number in the Virtual Chassis configuration file. Substitute the serial number of the replacement member switch (on the back of the switch) for the serial number of the member switch that was removed.

1. Power off and disconnect the member switch to be replaced.
2. If the replacement member switch has been previously configured, revert that switch's configuration to the factory defaults. See "Reverting to the Default Factory Configuration for the J-EX Series Switch" on page 341.
3. Connect and power on the replacement member switch.
4. Note the member ID displayed on the front panel.
5. Issue the **request virtual-chassis renumber** command from the Virtual Chassis master to change the member switch's current member ID to the member ID that belonged to the member switch that was removed from the Virtual Chassis configuration).

Remove a Member Switch and Make Its Member ID Available for Reassignment to a Different Switch

When you remove a member switch from the Virtual Chassis configuration, the master keeps its member ID on reserve. To make that member switch's member ID available for reassignment, issue the **request virtual-chassis recycle** command from the Virtual Chassis master.

NOTE: When you add or delete members in a Virtual Chassis configuration, internal routing changes might cause temporary traffic loss for a few seconds.

Related Documentation

- Monitoring Virtual Chassis Configuration Status and Statistics on page 809
- Adding a New Switch to an Existing Virtual Chassis Configuration (CLI Procedure) on page 786

Verifying That Graceful Routing Engine Switchover Is Working in the Virtual Chassis Configuration

Purpose Verify that graceful Routing Engine switchover (GRES) is working in the Virtual Chassis configuration.

Action On the master switch, verify the member ID of the backup Routing Engine:

```
{master:0}
user@switch> show virtual-chassis status
Virtual Chassis ID: 5efa.4b7a.aae6
                                                Mastership            Neighbor List
Member ID  Status   Serial No    Model    priority    Role      ID  Interface
0 (FPC 0)  Prsnt    BM0208105281 ex4200-24t      255  Master*    1  vcp-0
1 (FPC 1)  Prsnt    BP0208192350 ex4200-48t      255  Backup     0  vcp-0

Member ID for next new member: 2 (FPC 2)
```

Connect to the backup Routing Engine:

```
{master:0}
user@switch> request session member 1

{backup:1}
user@switch>
```

Verify that the backup Routing Engine is ready for switchover on member ID 1:

```
{backup:1}
user@switch> show system switchover

Graceful switchover: On
Configuration database: Ready
Kernel database: Ready
Peer state: Steady State
```

Switch the current backup Routing Engine to master Routing Engine:

NOTE: You must wait a minimum of 2 minutes between Routing Engine failovers for the Routing Engines to synchronize.

```
{backup:1}
user@switch> request chassis routing-engine master acquire
```

Verify that the master and backup Routing Engines have switched roles:

NOTE: Member ID 1 is now the master, and member ID 0 is now the backup.

```
{master:1}
user@switch> show virtual-chassis status

Virtual Chassis ID: 5efa.4b7a.aae6
```

```
                                              Mastership            Neighbor List
Member ID  Status   Serial No    Model    priority    Role      ID  Interface
0 (FPC 0)  Prsnt    BM0208105281 ex4200-24t      255  Backup     1  vcp-0
1 (FPC 1)  Prsnt    BP0208192350 ex4200-48t      255  Master*    0  vcp-0

Member ID for next new member: 2 (FPC 2)
```

Meaning With graceful Routing Engine switchover enabled, when you initiated a switchover from the backup Routing Engine, the backup Routing Engine became the master and the master Routing Engine became the backup.

Related Documentation

- Configuring Graceful Routing Engine Switchover in a Virtual Chassis (CLI Procedure) on page 801

CHAPTER 47

Troubleshooting Virtual Chassis

Troubleshooting a Virtual Chassis Configuration

Clear Virtual Chassis NotPrsnt Status and Make Member ID Available for Reassignment

Problem You disconnected a J-EX4200 from the Virtual Chassis configuration, but the disconnected switch's member ID is still displayed in the status output. You cannot reassign that member ID to another switch.

Solution When you disconnect a member of a Virtual Chassis configuration, the master retains the member ID and member configuration in its configuration database. The **show virtual-chassis status** command continues to display the member ID of the disconnected member with a status of **NotPrsnt**.

If want to permanently disconnect the member switch, you can free up the member ID by using the **request virtual-chassis recycle** command. This will also clear the status of that member.

Load Factory Default Does Not Commit on a Multimember Virtual Chassis

Problem The load factory default command fails on a multimember Virtual Chassis configuration.

Solution The **load factory default** command is not supported on a multimember Virtual Chassis configuration. For information on how to revert to factory default settings, see "Reverting to the Default Factory Configuration for the J-EX Series Switch" on page 341.

Member ID Persists When a Member Switch Is Disconnected From a Virtual Chassis

Problem Gigabit Ethernet interfaces retain their previous slot numbers when a member switch is disconnected from the Virtual Chassis configuration.

Solution

If a switch had been previously connected as a member of a Virtual Chassis configuration, it retains the member ID that it was assigned as a member of that configuration even after it is disconnected and operating as a standalone switch. The interfaces that were configured while the switch was a member of the Virtual Chassis configuration retain the old member ID as the first digit of the interface name.

For example, if the switch was previously member **1**, its interfaces are named **ge-1/0/0** and so on.

To change the switch's member ID, so that its member ID is **0**, and to rename the switch's interfaces accordingly, enter the following operational-mode commands:

1. To change the member ID to 0:

```
user@switch> request virtual-chassis renumber member-id 1 new-member-id 0
```

2. To rename the interfaces to match the new member ID:

```
user@switch# replace pattern ge-1/ with ge-0/
```

Related Documentation

- Monitoring Virtual Chassis Configuration Status and Statistics on page 809
- Configuring a Virtual Chassis (CLI Procedure) on page 781
- Configuring a Virtual Chassis (J-Web Procedure) on page 784
- For more information about the **replace** command, see *Junos OS CLI User Guide* at http://www.juniper.net/techpubs/software/junos/ .

CHAPTER 48

Configuration Statements for Virtual Chassis

[edit virtual-chassis] Configuration Statement Hierarchy

```
virtual-chassis {
  auto-sw-update {
    package-name-edit-virtual-chassis.xml package-name;
  }
  fast-failover (ge | vcp disable | xe);
  id id;
  mac-persistence-timer seconds;
  member member-id {
    mastership-priority number;
    no-management-vlan;
    serial-number;
    role;
  }
  no-split-detection;
  preprovisioned;
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable> <match
      regex>;
    flag flag ;
  }
}
```

Related Documentation

- Example: Configuring a Virtual Chassis with a Master and Backup in a Single Wiring Closet on page 717
- Example: Configuring a Virtual Chassis Interconnected Across Multiple Wiring Closets on page 733
- Example: Configuring a Virtual Chassis Using a Preprovisioned Configuration File on page 752
- Configuring a Virtual Chassis (CLI Procedure) on page 781
- Configuring a Virtual Chassis (J-Web Procedure) on page 784
- Virtual Chassis Overview on page 691

auto-sw-update

Syntax

```
auto-sw-update {
  package-name package-name;
}
```

Hierarchy Level [edit virtual-chassis]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable the automatic software update feature for Virtual Chassis configurations.

The remaining statement is explained separately.

Default The automatic software update feature is disabled.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Automatic Software Update on Virtual Chassis Member Switches on page 777
- Configuring Automatic Software Update on Virtual Chassis Member Switches (CLI Procedure) on page 800

fast-failover

Syntax fast-failover (ge | vcp disable | xe);

Hierarchy Level [edit virtual-chassis]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable the fast failover feature on all SFP uplink module Virtual Chassis ports (VCPs) or disable the fast failover feature on all dedicated VCPs in a ring topology.

Default Fast failover is enabled on dedicated VCPs; it is not enabled on uplink module VCPs.

Options

- **ge**—Enable fast failover on all Gigabit Ethernet uplink module VCPs in the ring.
- **vcp disable**—Disable fast failover on all dedicated VCPs in the ring.
- **xe**—Enable fast failover on all 10-Gigabit Ethernet uplink module VCPs in the ring.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Fast Failover on Uplink Module VCPs to Reroute Traffic When a Virtual Chassis Member Switch or Inter-Member Link Fails on page 763
- Configuring Fast Failover in a Virtual Chassis Configuration on page 798
- Disabling Fast Failover in a Virtual Chassis Configuration on page 799

graceful-switchover

Syntax graceful-switchover;

Hierarchy Level [edit **chassis redundancy**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For switches with more than one Routing Engine, including those in a Virtual Chassis, configure the master Routing Engine to switch over gracefully to a backup Routing Engine without interruption to packet forwarding.

Default Graceful Routing Engine switchover (GRES) is disabled.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Nonstop Active Routing on J-EX Series Switches
- Configuring Graceful Routing Engine Switchover in a J-EX4200 Virtual Chassis (CLI Procedure) on page 801
- Configuring Nonstop Active Routing on J-EX Series Switches (CLI Procedure)
- Installing Software on a J-EX8200 Switch with Redundant Routing Engines (CLI Procedure) on page 71

id

Syntax id *id*;

Hierarchy Level [edit **virtual-chassis**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the alphanumeric string that identifies a Virtual Chassis configuration.

Options *id*—ID of the Virtual Chassis configuration, which uses the ISO family address format—for example, **9622.6ac8.5345**.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge on page 767
- Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge (CLI Procedure) on page 800
- Understanding Split and Merge in a Virtual Chassis Configuration on page 712

mac-persistence-timer

Syntax mac-persistence-timer *minutes*;

Hierarchy Level [edit virtual-chassis]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description If the master is physically disconnected or removed from the Virtual Chassis configuration, the MAC persistence timer determines how long the backup (new master) continues to use the address of the old master. When the MAC persistence timer expires, the backup (new master) begins to use its own MAC address.

There are no minimum or maximum timer limits.

Default 10 minutes

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring the Timer for the Backup Member to Start Using Its Own MAC Address, as Master of Virtual Chassis (CLI Procedure) on page 797
- Understanding Virtual Chassis Components on page 694

mastership-priority

Syntax mastership-priority *number* ;

Hierarchy Level [edit **virtual-chassis member** *member-id*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description The mastership priority value is the most important factor in determining the role of the J-EX4200 member switch within the Virtual Chassis configuration. Other factors (see "Understanding How the Master in a Virtual Chassis Configuration Is Elected" on page 698) also affect the election of the master.

The mastership priority value takes the highest precedence in the master election algorithm. The member switch with highest mastership priority becomes the master of the Virtual Chassis configuration. Toggling back and forth between master and backup status in failover conditions is undesirable, so we recommend that you assign the same mastership priority value to both the master and the backup. Secondary factors in the master election algorithm determine which of these two members (that is, the two members that are assigned the highest mastership priority value) functions as the master of the Virtual Chassis configuration.

Default 128

Options *number*—Mastership priority value.
Range: 1 through 255

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring a Virtual Chassis with a Master and Backup in a Single Wiring Closet on page 717
- Example: Configuring a Virtual Chassis Interconnected Across Multiple Wiring Closets on page 733
- Configuring a Virtual Chassis (CLI Procedure) on page 781
- Configuring a Virtual Chassis (J-Web Procedure) on page 784
- Understanding Virtual Chassis Components on page 694

member

Syntax

```
member member-id {
  mastership-priority number;
  no-management-vlan;
  serial-number;
  role;
}
```

Hierarchy Level [edit virtual-chassis]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a J-EX4200 switch as a member of a Virtual Chassis configuration.

Default When a J-EX4200 is powered on as a standalone switch (not interconnected through its Virtual Chassis ports with other J-EX4200 switches), its default member ID is **0**.

Options *member-id*—Identifies a specific member switch of a Virtual Chassis configuration.

Range: 0 through 9

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring a Virtual Chassis Using a Preprovisioned Configuration File on page 752
- Configuring a Virtual Chassis (CLI Procedure) on page 781
- Configuring a Virtual Chassis (J-Web Procedure) on page 784
- Understanding Virtual Chassis Components on page 694

no-management-vlan

Syntax no-management-vlan;

Hierarchy Level [edit **virtual-chassis member** *member-id*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Remove the specified member's out-of-band management port from the Virtual Management Ethernet (VME) global management VLAN of the Virtual Chassis configuration.

For a member that is functioning in a linecard role, you can use this configuration to reserve the member's management Ethernet port for local troubleshooting:

```
virtual-chassis {
  member 2 {
    no-management-vlan;
    }
  }
```

You cannot configure the IP address for a local management Ethernet port using the CLI or the J-Web interface. To do this, you need to use the shell **ifconfig** command.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Setting Up a Multimember Virtual Chassis Access Switch with a Default Configuration on page 727
- Configuring the Virtual Management Ethernet Interface for Global Management of a Virtual Chassis (CLI Procedure) on page 797
- Understanding Global Management of a Virtual Chassis Configuration on page 699

no-split-detection

Syntax no-split-detection;

Hierarchy Level [edit virtual-chassis]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable the split and merge feature in a Virtual Chassis configuration. The split and merge feature is enabled by default on J-EX4200 switches.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge on page 767
- Disabling Split and Merge in a Virtual Chassis Configuration (CLI Procedure) on page 799
- Assigning the Virtual Chassis ID to Determine Precedence During a Virtual Chassis Merge (CLI Procedure) on page 800
- Understanding Split and Merge in a Virtual Chassis Configuration on page 712

package-name

Syntax package-name *package-name*;

Hierarchy Level [edit virtual-chassis **auto-sw-update**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the software package name or location of the software package to be used by the automatic software update feature for Virtual Chassis configurations.

Default No package name is specified.

Options ***package-name***—Name of the software package or the URL to the software package to be used.

- If the software package is located on a local directory on the switch, use the following format for ***package-name***:

  /pathname/package-name

- If the software package is to be downloaded and installed from a remote location, use one of the following formats:

  ftp:*//hostname/pathname/package-name*

  ftp:*//username*:*prompt*@ftp.*hostname.net/package-name*

  http:*//hostname/pathname/package-name*

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Automatic Software Update on Virtual Chassis Member Switches on page 777
- Configuring Automatic Software Update on Virtual Chassis Member Switches (CLI Procedure) on page 800

preprovisioned

Syntax preprovisioned;

Hierarchy Level [edit virtual-chassis]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable the preprovisioned configuration mode for a Virtual Chassis configuration.

When preprovisioned configuration mode is enabled, you cannot use the CLI or the J-Web interface to change the mastership priority or member ID of member switches.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring a Virtual Chassis Using a Preprovisioned Configuration File on page 752
- Configuring a Virtual Chassis (CLI Procedure) on page 781
- Configuring a Virtual Chassis (J-Web Procedure) on page 784
- Adding a New Switch to an Existing Virtual Chassis Configuration (CLI Procedure) on page 786
- Replacing a Member Switch of a Virtual Chassis Configuration (CLI Procedure) on page 811
- Understanding Virtual Chassis Configuration on page 704

redundancy (Graceful Switchover)

Syntax

```
redundancy {
  graceful-switchover;
}
```

Hierarchy Level [edit chassis]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For J-EX4200 switches configured as a Virtual Chassis and for J-EX8200 switches with more than one Routing Engine, enable redundant Routing Engines.

The remaining statement is explained separately.

Default Redundancy is enabled for the Routing Engines.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Graceful Routing Engine Switchover in a J-EX4200 Virtual Chassis (CLI Procedure) on page 801
- Installing Software on a J-EX8200 Switch with Redundant Routing Engines (CLI Procedure) on page 71

role

Syntax role (routing-engine | line-card);

Hierarchy Level [edit **virtual-chassis preprovisioned member** *member-id*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description In a preprovisioned Virtual Chassis configuration, specify the role to be performed by each J-EX4200 member switch. Associate the role permanently with the member's serial number.

Options **routing-engine**—Enables the member eligible to function as a master or backup of the Virtual Chassis configuration. The master manages all the members of the Virtual Chassis configuration and runs the chassis management processes and control protocols. The backup synchronizes with the master in terms of protocol states, forwarding tables, and so forth, so that it is prepared to preserve routing information and maintain network connectivity without disruption in case the master is unavailable.

Specify two and only two members as routing-engine. The software determines which of the two members assigned the routing-engine role functions as master, based on the master election algorithm. See "Understanding How the Master in a Virtual Chassis Configuration Is Elected" on page 698.

line-card—Enables the member to be eligible to function only in the linecard role. Any member of the Virtual Chassis configuration other than the master or backup functions in the linecard role and runs only a subset of the Junos OS for J-EX Series switches. A member functioning in the linecard role does not run the chassis control protocols. A Virtual Chassis configuration must have at least three members in order to include a member that functions in the linecard role.

When you use a preprovisioned configuration, you cannot modify the mastership priority or member ID of member switches through the user interfaces. The mastership priority value is generated by the software, based on the assigned role:

- A member configured as **routing-engine** is assigned the mastership priority **129**.
- A member configured as **line-card** is assigned the mastership priority **0**.
- A member listed in the preprovisioned configuration without an explicitly specified role is assigned the mastership priority **128**.

The configured role specifications are permanent. If both **routing-engine** members should fail, a **line-card** member cannot take over as master of the Virtual Chassis configuration. You must delete the preprovisioned configuration in order to change the specified roles.

It is possible to explicitly configure two members as **routing-engine** and to configure additional switches as members of the preprovisioned Virtual Chassis by specifying only their serial numbers. If you do not explicitly configure the role of the additional

members, they function in a linecard role by default. In that case, a member that is functioning in a linecard role can take over mastership if the members functioning as master and backup (routing-engine role) both fail.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring a Virtual Chassis Using a Preprovisioned Configuration File on page 752
- Configuring a Virtual Chassis (CLI Procedure) on page 781
- Configuring a Virtual Chassis (J-Web Procedure) on page 784
- Adding a New Switch to an Existing Virtual Chassis Configuration (CLI Procedure) on page 786
- Replacing a Member Switch of a Virtual Chassis Configuration (CLI Procedure) on page 811
- Understanding Virtual Chassis Configuration on page 704

serial-number

Syntax serial-number *serial-number*;

Hierarchy Level [edit **virtual-chassis preprovisioned member** *member-id*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description In a preprovisioned Virtual Chassis configuration, specify the serial number of each J-EX4200 member switch to be included in the Virtual Chassis configuration. If you do not include the serial number within the Virtual Chassis configuration, the switch cannot be recognized as a member of a preprovisioned configuration.

Options **serial-number**—The switch's permanent serial number, which is located on the back of the switch.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring a Virtual Chassis Using a Preprovisioned Configuration File on page 752
- Configuring a Virtual Chassis (CLI Procedure) on page 781
- Configuring a Virtual Chassis (J-Web Procedure) on page 784
- Adding a New Switch to an Existing Virtual Chassis Configuration (CLI Procedure) on page 786
- Replacing a Member Switch of a Virtual Chassis Configuration (CLI Procedure) on page 811
- Understanding Virtual Chassis Configuration on page 704

traceoptions

Syntax

```
traceoptions {
  file filename <files number> <no-stamp> <replace> <size size> <world-readable |
    no-world-readable>;
  flag flag <detail> <disable> <receive> <send>;
}
```

Hierarchy Level [edit virtual-chassis]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define tracing operations for the Virtual Chassis configuration.

Default Tracing operations are disabled.

Options **detail**—(Optional) Generate detailed trace information for a flag.

disable—(Optional) Disable a flag.

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000

Default: 3 files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:

- **all**—All tracing operations.

TIP: The all flag displays a subset of logs that are useful in debugging most issues. For more detailed information, use all detail.

- **auto-configuration**—Trace Virtual Chassis ports (VCPs) that have been automatically configured.
- **csn**—Trace Virtual Chassis complete sequence number (CSN) packets.
- **error**—Trace Virtual Chassis errored packets.
- **hello**—Trace Virtual Chassis hello packets.
- **krt**—Trace Virtual Chassis KRT events.
- **lsp**—Trace Virtual Chassis link-state packets.

- **lsp-generation**—Trace Virtual Chassis link-state packet generation.
- **me**—Trace Virtual Chassis ME events.
- **normal**—Trace normal events.
- **packets**—Trace Virtual Chassis packets.
- **parse**—Trace reading of the configuration.
- **psn**—Trace partial sequence number (PSN) packets.
- **route**—Trace Virtual Chassis routing information.
- **spf**—Trace Virtual Chassis SPF events.
- **state**—Trace Virtual Chassis state transitions.
- **task**—Trace Virtual Chassis task operations.

no-stamp—(Optional) Do not place a timestamp on any trace file.

no-world-readable—(Optional) Restrict file access to the user who created the file.

receive—(Optional) Trace received packets.

replace—(Optional) Replace a trace file rather than appending information to it.

send—(Optional) Trace transmitted packets.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **files** option.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10 KB through 1 GB

Default: 128 KB

world-readable—(Optional) Enable unrestricted file access.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Monitoring Virtual Chassis Configuration Status and Statistics on page 809
- Verifying the Member ID, Role, and Neighbor Member Connections of a Virtual Chassis Member on page 807
- Verifying That the Virtual Chassis Ports Are Operational on page 808
- Troubleshooting a Virtual Chassis Configuration on page 815

virtual-chassis

Syntax

```
virtual-chassis {
  auto-sw-upgrade {
    package-name-edit-virtual-chassis.xml package-name;
  }
  fast-failover (ge | vcp disable | xe);
  id id;
  mac-persistence-timer seconds;
  member member-id {
    mastership-priority number;
    no-management-vlan;
    serial-number;
    role;
  }
  no-split-detection;
  preprovisioned;
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable> <match
      regex>;
    flag flag ;
  }
}
```

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure Virtual Chassis information on a J-EX4200 switch.

The remaining statements are explained separately.

Default A standalone J-EX4200 switch is a Virtual Chassis by default. It has a default member ID of 0, a default mastership priority of 128, and a default role as master.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring a Virtual Chassis with a Master and Backup in a Single Wiring Closet on page 717
- Configuring a Virtual Chassis (CLI Procedure) on page 781
- Configuring a Virtual Chassis (J-Web Procedure) on page 784
- Understanding Virtual Chassis Components on page 694

CHAPTER 49

Operational Mode Commands for Virtual Chassis

clear virtual-chassis vc-port statistics

Syntax

clear virtual-chassis vc-port statistics
<all-members>
<*interface-name*>
<local>
<member *member-id*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Clear—reset to zero (0)—the traffic statistics counters on Virtual Chassis ports (VCPs).

Options

none—Clear traffic statistics for the VCPs of all members of a Virtual Chassis configuration.

all-members—(Optional) Clear traffic statistics for the VCPs of all members of a Virtual Chassis configuration.

interface-name—(Optional) Name of the VCP interface to be cleared of its traffic statistics. Specify either **vcp-0** or **vcp-1**.

local—(Optional) Clear VCP traffic statistics from only the switch on which this command is entered.

member *member-id*—(Optional) Clear VCP traffic statistics from only the specified member of a Virtual Chassis configuration.

Required Privilege Level

clear

Related Documentation

- **show virtual-chassis vc-port statistics on page 854**
- **show virtual-chassis vc-port on page 851**
- Monitoring Virtual Chassis Configuration Status and Statistics on page 809
- Understanding Virtual Chassis Components on page 694

List of Sample Output

clear virtual-chassis vc-port statistics on page 836
clear virtual-chassis vc-port statistics member 3 on page 836

clear virtual-chassis vc-port statistics

```
user@SWA-0> clear virtual-chassis vc-port statistics
fpc0:
--------------------------------------------------------------------------
Statistics cleared

{master:0}
```

clear virtual-chassis vc-port statistics member 3

```
user@SWA-0> clear virtual-chassis vc-port statistics member 3
Cleared statistics on member 3
```

request session member

Syntax **request session member *member-id***

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Starts a session with the specified member of a Virtual Chassis configuration.

Options ***member-id***—Select the specific member of the Virtual Chassis configuration with which you want to establish a session.

Required Privilege Level maintenance

Related Documentation

- **member on page 823**
- Understanding Virtual Chassis Components on page 694

request virtual-chassis recycle

Syntax **request virtual-chassis recycle member-id *member-id***

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Make a previously used member ID available for reassignment.

When you remove a member switch from the Virtual Chassis configuration, the master reserves that member ID. To make the member ID available for reassignment, you must use this command.

NOTE: You can run this command from the Virtual Chassis master only.

Options member-id *member-id*—Specify the member id that you want to make available for reassignment to a different member switch.

Required Privilege Level system-control

Related Documentation

- **request virtual-chassis renumber on page 839**
- Replacing a Member Switch of a Virtual Chassis Configuration (CLI Procedure) on page 811

List of Sample Output **request virtual-chassis recycle member-id 3 on page 838**

request virtual-chassis recycle member-id 3

```
user@host> request virtual-chassis recycle member-id 3
```

request virtual-chassis renumber

Syntax request virtual-chassis renumber member-id *old-member-id* new-member-id *new-member-id*

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Renumber a member of a Virtual Chassis configuration.

NOTE: You can run this command from the Virtual Chassis master only.

Options member-id *old-member-id*—Specify the ID of the member that you wish to renumber.

new-member-id *new-member-id*—Specify an unassigned member ID (from 0 through 9).

Required Privilege Level system-control

Related Documentation

- **request virtual-chassis recycle on page 838**
- Replacing a Member Switch of a Virtual Chassis Configuration (CLI Procedure) on page 811

List of Sample Output **request virtual-chassis renumber member-id 5 new-member-id 4 on page 839**

request virtual-chassis renumber member-id 5 new-member-id 4

```
user@SWA-0> request virtual-chassis renumber member-id 5 new-member-id 4
```

request virtual-chassis vc-port

Syntax **request virtual-chassis vc-port set | delete pic-slot** *pic-slot* **port** *port-number* **<member** *member-id***>**

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable or disable an uplink module port (on an SFP or SFP+ uplink module) or an SFP network port on a J-EX4200-24F switch as a Virtual Chassis port (VCP).

Options

pic-slot *pic-slot*—Number of the PIC slot for the uplink module port or SFP network port on a J-EX4200-24F switch. Specify **1** to represent the uplink module PIC on the J-EX Series switch. Specify **0** to represent the SFP network port on a J-EX4200-24F switch.

port *port-number*—Number of the uplink module port (**0** or **1**) or SFP network port on a J-EX4200-24F switch (**0** to **23**) that is to be enabled or disabled as a VCP.

member *member-id*—(Optional) Enable or disable the specified VCP on the specified member of the Virtual Chassis configuration.

Additional Information If you omit **member** ***member-id***, this command defaults to enabling or disabling the uplink VCP or SFP network port configured as a VCP on the switch where the command is issued.

Required Privilege Level system-control

Related Documentation

- **request virtual-chassis vc-port on page 841** (dedicated port)
- **show virtual-chassis vc-port on page 851**
- **show virtual-chassis vc-port statistics on page 854**
- **clear virtual-chassis vc-port statistics on page 836**
- Understanding Virtual Chassis Components on page 694

List of Sample Output

request virtual-chassis vc-port set pic-slot 1 port 0 on page 840
request virtual-chassis vc-port set pic-slot 1 port 1 member 3 on page 840
request virtual-chassis vc-port delete pic-slot 1 port 1 member 3 on page 840

request virtual-chassis vc-port set pic-slot 1 port 0

```
user@host>request virtual-chassis vc-port set pic-slot 1 port 0
```

To check the results of this command, use the **show virtual-chassis vc-port** command.

request virtual-chassis vc-port set pic-slot 1 port 1 member 3

```
user@host>request virtual-chassis vc-port set pic-slot 1 port 1 member 3
```

To check the results of this command, use the **show virtual-chassis vc-port** command.

request virtual-chassis vc-port delete pic-slot 1 port 1 member 3

```
user@host>request virtual-chassis vc-port delete pic-slot 1 port 1 member 3
```

To check the results of this command, use the **show virtual-chassis vc-port** command.

request virtual-chassis vc-port

Syntax **request virtual-chassis vc-port set interface *vcp-interface-name*** **<member *member-id*> <disable>**

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable or enable a Virtual Chassis port (VCP) for a dedicated VCP on the rear panel of the Virtual Chassis.

Options

interface *vcp-interface-name* —Name of the interface to enable or disable. Specify either **vcp-0** or **vcp-1**.

member *member-id* —(Optional) Enable or disable the specified VCP on the specified member of the Virtual Chassis configuration.

disable —(Optional) Disable the specified VCP. If you omit this keyword, the command enables the dedicated VCP.

Additional Information If you omit **member *member-id***, this command defaults to disabling or enabling the dedicated VCP on the switch where the command is issued. The dedicated VCPs are enabled in the factory default configuration.

Required Privilege Level system-control

Related Documentation

- **request virtual-chassis vc-port on page 840**
- **show virtual-chassis vc-port on page 851**
- **show virtual-chassis vc-port statistics on page 854**
- **clear virtual-chassis vc-port statistics on page 836**
- Understanding Virtual Chassis Components on page 694

List of Sample Output

request virtual-chassis vc-port set interface vcp-0 disable on page 841
request virtual-chassis vc-port set interface vcp-0 member 3 disable on page 841

request virtual-chassis vc-port set interface vcp-0 disable

```
user@host> request virtual-chassis vc-port set interface vcp-0 disable
```

To check the results of this command, use the **show virtual-chassis vc-port** command.

request virtual-chassis vc-port set interface vcp-0 member 3 disable

```
user@host> request virtual-chassis vc-port set interface vcp-0 member 3 disable
```

To check the results of this command, use the **show virtual-chassis vc-port** command.

show system uptime

Syntax **show system uptime (all-members | member *member-id)***

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display the current time and information about how long the Virtual Chassis, Virtual Chassis software, and routing protocols have been running.

Options

all-members—Display the current time and information about how long the Virtual Chassis, Virtual Chassis software, and routing protocols have been running for all the member switches of the Virtual Chassis configuration.

member *member-id*—Display the current time and information about how long the Virtual Chassis, Virtual Chassis software, and routing protocols have been running for the specific member of the Virtual Chassis configuration.

Required Privilege Level view

Related Documentation

- **virtual-chassis on page 834**
- Monitoring System Properties on page 550
- For more information about **show system uptime**, see the *Junos OS System Basics Services and Command Reference* at http://www.juniper.net/techpubs/software/junos/.

List of Sample Output **show system uptime member 0 on page 843**

Output Fields Table 110 on page 842 lists the output fields for the **show system uptime** command. Output fields are listed in the approximate order in which they appear.

Table 110: show system uptime Output Fields

Field Name	Field Description
Current time	Current system time in UTC.
System booted	Date and time when the switch was last booted and how long it has been running.
Protocols started	Date and time when the routing protocols were last started and how long they have been running.
Last configured	Date and time when a configuration was last committed. Also shows the name of the user who issued the last commit command.
Time and up	Current time, in the local time zone, and how long the switch has been operational.
Users	Number of users logged into the switch.

Table 110: show system uptime Output Fields *(continued)*

Field Name	Field Description
Load averages	Load averages for the last 1 minute, 5 minutes, and 15 minutes.

show system uptime member 0

```
user@host>show system uptime member 0
fpc0:
--------------------------------------------------------------------------
Current time: 2008-02-06 05:24:20 UTC
System booted: 2008-01-31 08:26:54 UTC (5d 20:57 ago)
Protocols started: 2008-01-31 08:27:56 UTC (5d 20:56 ago)
Last configured: 2008-02-05 03:26:43 UTC (1d 01:57 ago) by root
 5:24AM  up 5 days, 20:57, 1 user, load averages: 0.14, 0.06, 0.01
```

show virtual-chassis active topology

Syntax **show virtual-chassis active-topology**
<(all-members | member *member-id*)>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display the active topology of the Virtual Chassis configuration with reachability information.

Options none—Display the active topology of the member switch where the command is issued.

all-members—Display the active topology of all members of the Virtual Chassis configuration.

member *member-id*—Display the active topology of a specified member of the Virtual Chassis configuration.

Required Privilege Level view

Related Documentation
- Monitoring Virtual Chassis Configuration Status and Statistics on page 809
- Understanding Virtual Chassis Configuration on page 704

List of Sample Output **show virtual-chassis active-topology on page 844**

Output Fields Table 111 on page 844 lists the output fields for the **show virtual-chassis active-topology** command. Output fields are listed in the approximate order in which they appear.

Table 111: show virtual-chassis active-topology Output Fields

Field Name	Field Description
Destination ID	Specifies the member ID of the destination.
Next-hop	Specifies the member ID and VCP of the next-hop to which packets for the destination ID are forwarded.

show virtual-chassis active-topology

```
user@SWA-0> show virtual-chassis active-topology
  1                   1(vcp-1)

  2                   1(vcp-1)

  3                   1(vcp-1)

  4                   1(vcp-1)
```

```
5                   8(vcp-0)  1(vcp-1)

6                   8(vcp-0)

7                   8(vcp-0)

8                   8(vcp-0)
```

show virtual-chassis fast-failover

Syntax **show virtual-chassis fast-failover**

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display information about the fast failover feature in a Virtual Chassis configuration.

Required Privilege Level view

Related Documentation

- Example: Configuring Fast Failover on Uplink Module VCPs to Reroute Traffic When a Virtual Chassis Member Switch or Inter-Member Link Fails on page 763
- Configuring Fast Failover in a Virtual Chassis Configuration on page 798
- Disabling Fast Failover in a Virtual Chassis Configuration on page 799
- Understanding Fast Failover in a Virtual Chassis Configuration on page 706

List of Sample Output **show virtual-chassis fast-failover on page 846**

Output Fields Table 112 on page 846 lists the output fields for the **show virtual-chassis fast-failover** command. Output fields are listed in the approximate order in which they appear.

Table 112: show virtual-chassis fast-failover Output Fields

Field Name	Field Description
Fast failover on dedicated VCP ports	Indicates fast failover status on dedicated VCPs.
Fast failover on XE uplink VCP ports	Indicates fast failover status on XFP uplink module VCPs. (Not supported on Dell PowerConnect J-EX Series Switches.)
Fast failover on GE uplink VCP ports	Indicates fast failover status on SFP uplink module VCPs.

show virtual-chassis fast-failover

```
user@switch1> show virtual-chassis fast-failover
Fast failover on dedicated VCP ports: Enabled
Fast failover on XE uplink VCP ports: Disabled
Fast failover on GE uplink VCP ports: Enabled
```

show virtual-chassis status

Syntax **show virtual-chassis status**

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display information about all the members of the Virtual Chassis configuration.

Options none—Display all information for all member switches of the Virtual Chassis configuration.

Required Privilege Level view

Related Documentation
- Monitoring Virtual Chassis Configuration Status and Statistics on page 809
- Understanding Virtual Chassis Configuration on page 704

Output Fields Table 113 on page 847 lists the output fields for the **show virtual-chassis status** command. Output fields are listed in the approximate order in which they appear.

Table 113: show virtual-chassis status Output Fields

Field Name	Field Description
Virtual Chassis ID	Assigned ID that applies to the entire Virtual Chassis configuration.
Member ID	Assigned member ID and FPC slot (from 0 through 9).
Status	For a nonprovisioned configuration: • **Prsnt** for a member that is currently connected to the Virtual Chassis configuration • **NotPrsnt** for a member ID that has been assigned but is not currently connected For a preprovisioned configuration: • **Prsnt** for a member that is specified in the preprovisioned configuration file and is currently connected to the Virtual Chassis configuration. • **Unprvsnd** for a member that is interconnected with the Virtual Chassis configuration, but is not specified in the preprovisioned configuration file.
Serial No	Serial number of the member switch.
Model	Model number of the member switch.
Mastership Priority	Mastership priority value of the member switch.
Role	Role of the member switch.
Neighbor List	Member ID of the neighbor member to which this member's VCP interface is connected.

show virtual-chassis status

```
user@SWA-0> show virtual-chassis status

   Virtual Chassis ID: 0019.e250.47a0
                                                 Mastership            Neighbor List
Member ID  Status   Serial No    Model    priority    Role      ID  Interface
0 (FPC 0)  Prsnt    AK0207360276 ex4200-24t      249  Master*    8  vcp-0
                                                                1  vcp-1
1 (FPC 1)  Prsnt    AK0207360281 ex4200-24t      248  Backup     0  vcp-0
                                                                2  vcp-1
2 (FPC 2)  Prsnt    AJ0207391130 ex4200-48t      247  Linecard   1  vcp-0
                                                                3  vcp-1
3 (FPC 3)  Prsnt    AK0207360280 ex4200-24t      246  Linecard   2  vcp-0
                                                                4  vcp-1
4 (FPC 4)  Prsnt    AJ0207391113 ex4200-48t      245  Linecard   3  vcp-0
                                                                5  vcp-1
5 (FPC 5)  Prsnt    BP0207452204 ex4200-48t      244  Linecard   4  vcp-0
                                                                6  vcp-1
6 (FPC 6)  Prsnt    BP0207452222 ex4200-48t      243  Linecard   5  vcp-0
                                                                7  vcp-1
7 (FPC 7)  Prsnt    BR0207432028 ex4200-24f      242  Linecard   6  vcp-0
                                                                8  vcp-1
8 (FPC 8)  Prsnt    BR0207431996 ex4200-24f      241  Linecard   7  vcp-0
                                                                0  vcp-1

Member ID for next new member: 9 (FPC 9)
```

show virtual-chassis vc-path

Syntax **show virtual-chassis vc-path source-interface *interface-name* destination-interface *interface-name***

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Show the path a packet takes when going from a source interface to a destination interface in a Virtual Chassis configuration.

Options source-interface *interface-name* —Name of the interface from which the packet originates

destination-interface *interface-name* —Name of the interface to which the packet is delivered

Required Privilege Level view

Related Documentation
- Monitoring Virtual Chassis Configuration Status and Statistics on page 809
- Understanding Virtual Chassis Configuration on page 704

List of Sample Output **show virtual-chassis vc-path source-interface destination-interface on page 849**

Output Fields Table 114 on page 849 lists the output fields for the **show virtual-chassis vc-path** command. Output fields are listed in the approximate order in which they appear.

Table 114: show virtual-chassis vc-path Output Fields

Field Name	Field Description
Hop	The number of hops between the source and destination interfaces.
Member	The Virtual Chassis ID of the member switch that contains the Packet Forwarding Engine for each intermediate hop.
PFE-Device	The number of the Packet Forwarding Engine in each Virtual Chassis member through which a packet passes. Each Packet Forwarding Engine is the next hop of the preceding Packet Forwarding Engine.
Interface	The name of the interface through which the Packet Forwarding Engines are connected. The interface for the first hop is always the source interface and the interface for the last hop is always the destination interface. For intermediate hops, the **Interface** field denotes the Packet Forwarding Engines through which the packet passes on its way to the next hop.

show virtual-chassis vc-path source-interface destination-interface

```
user@switch> show virtual-chassis vc-path source-interface ge-0/0/0 destination-interface
ge-1/0/1
vc-path from ge-0/0/0 to ge-1/0/1
Hop      Member     PFE-Device      Interface
0         0          1              ge-0/0/0
1         0          0              internal-1/24
```

```
2          1          3                vcp-0
3          1          4                ge-1/0/1
```

show virtual-chassis vc-port

Syntax

show virtual-chassis vc-port
<(all-members | member *member-id*)>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display the status of the Virtual Chassis ports (VCPs), including both the dedicated VCPs and the uplink module ports configured as VCPs.

Options

none—Display the operational status of all the VCPs of the member switch where the command is issued.

all-members—(Optional) Display the operational status of all the VCPs on all members of the Virtual Chassis configuration.

member *member-id*—(Optional) Display the operational status of all the VCPs for the specified member of the Virtual Chassis configuration.

Required Privilege Level

view

Related Documentation

- **show virtual-chassis vc-port statistics on page 854**
- Monitoring Virtual Chassis Configuration Status and Statistics on page 809
- Understanding Virtual Chassis Configuration on page 704

List of Sample Output

show virtual-chassis vc-port on page 852
show virtual-chassis vc-port all-members on page 852

Output Fields

Table 115 on page 851 lists the output fields for the **show virtual-chassis vc-port** command. Output fields are listed in the approximate order in which they appear.

Table 115: show virtual-chassis vc-port Output Fields

Field Name	Field Description
fpc*number*	The FPC number is the same as the member ID.
Interface or PIC/Port	VCP interface name. Unlike network interface names, a VCP interface name does not include a slot number (member ID). • The dedicated VCPs are **vcp-0** and **vcp-1**. • The uplink module ports set as VCPs are named **1/0** and **1/1**, representing the PIC number and the port number.

Table 115: show virtual-chassis vc-port Output Fields *(continued)*

Field Name	Field Description
Type	Type of VCP: • **Dedicated** (on the rear panel) • **Configured** (uplink module port configured as a VCP) • **Auto-Configured** (uplink module port autoconfigured as a VCP) See "Setting an Uplink Module Port as a Virtual Chassis Port (CLI Procedure)" on page 792 for information about configuring VCPs.
Trunk ID	A positive-number ID assigned to a LAG formed by the Virtual Chassis. The trunk ID value is **–1** if no trunk is formed. A LAG between uplink VCPs requires that the link speed be the same on connected interfaces and that at least two VCPs on one member be connected to at least two VCPs on the other member. Dedicated VCP LAGs are assigned trunk IDs 1 and 2. Trunk IDs for LAGs formed with uplink VCPs therefore have values of 3 or greater. The trunk ID value changes if the link-adjacency state between LAG members changes; trunk membership is then allocated or deallocated.
Status	Interface status: **down** or **up**.
Speed (mbps)	Speed of the interface in megabits per second.
Neighbor ID/Interface	The Virtual Chassis member ID and interface of a VCP on a member switch that is connected to the interface or PIC/Port field in the same row as this interface.

show virtual-chassis vc-port

```
user@switch> show virtual-chassis vc-port

fpc0:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status       Speed        Neighbor
or                             ID                (mbps)       ID  Interface
PIC / Port
vcp-0       Dedicated           1    Up           32000        1   vcp-1
vcp-1       Dedicated           2    Up           32000        0   vcp-0
1/0         Auto-Configured     3    Up           1000         2   vcp-255/1/0
1/0         Auto-Configured     3    Up           1000         2   vcp-255/1/1
```

show virtual-chassis vc-port all-members

```
user@switch> show virtual-chassis vc-port all-members

fpc0:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status       Speed        Neighbor
or                             ID                (mbps)       ID  Interface
PIC / Port
vcp-0       Dedicated           1    Up           32000        1   vcp-1
vcp-1       Dedicated           2    Up           32000        0   vcp-0
1/0         Auto-Configured     3    Up           1000         2   vcp-255/1/0
1/1         Auto-Configured     3    Up           1000         2   vcp-255/1/1

fpc1:
```

```
--------------------------------------------------------------------------
Interface   Type              Trunk  Status       Speed        Neighbor
or                             ID                (mbps)       ID  Interface
PIC / Port
vcp-0       Dedicated           1    Up           32000        0   vcp-1
vcp-1       Dedicated           2    Up           32000        0   vcp-0
1/0         Auto-Configured    -1    Up           1000         3   vcp-255/1/0

fpc2:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status       Speed        Neighbor
or                             ID                (mbps)       ID  Interface
PIC / Port
vcp-0       Dedicated           1    Up           32000        3   vcp-1
vcp-1       Dedicated           2    Up           32000        3   vcp-0
1/0         Auto-Configured     3    Up           1000         0   vcp-255/1/0
1/1         Auto-Configured     3    Up           1000         0   vcp-255/1/1

fpc3:
--------------------------------------------------------------------------
Interface   Type              Trunk  Status       Speed        Neighbor
or                             ID                (mbps)       ID  Interface
PIC / Port
vcp-0       Dedicated           1    Up           32000        2   vcp-0
vcp-1       Dedicated           2    Up           32000        2   vcp-1
1/0         Auto-Configured    -1    Up           1000         1   vcp-255/1/0
```

show virtual-chassis vc-port statistics

Syntax

show virtual-chassis vc-port statistics
<all-members>
<brief | detail | extensive >
<*interface-name*>
<local>
<member *member-id*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display the traffic statistics collected on Virtual Chassis ports (VCPs).

Options

none—Display traffic statistics for the VCPs of all members of a Virtual Chassis configuration.

brief | detail | extensive—(Optional) Display the specified level of output. Using the **brief** option is equivalent to entering the command with no options (the default). The **detail** and **extensive** options provide identical displays.

all-members—(Optional) Display traffic statistics for the VCPs of all members of a Virtual Chassis configuration.

interface-name—(Optional) Name of the VCP interface for which to display traffic statistics. Specify either **vcp-0** or **vcp-1** or an internal port in the VCP subsystem—for example, **internal-0/24**.

local—(Optional) Display VCP traffic statistics for only the switch on which this command is entered.

member *member-id*—(Optional) Display VCP traffic statistics for only the specified member of a Virtual Chassis configuration.

Required Privilege Level

view

Related Documentation

- **clear virtual-chassis vc-port statistics on page 836**
- **show virtual-chassis vc-port on page 851**
- Monitoring Virtual Chassis Configuration Status and Statistics on page 809

List of Sample Output

show virtual-chassis vc-port statistics on page 857
show virtual-chassis vc-port statistics brief on page 857
show virtual-chassis vc-port statistics extensive on page 857
show virtual-chassis vc-port statistics member 0 on page 858

Output Fields

Table 116 on page 855 lists the output fields for the **show virtual-chassis vc-port statistics** command. Output fields are listed in the approximate order in which they appear.

Table 116: show virtual-chassis vc-port statistics Output Fields

Field Name	Field Description	Level of Output
fpc*number*	ID of the Virtual Chassis member. The FPC number is the same as the member ID.	All levels
Interface	VCP interface name. Unlike network interface names, a VCP interface does not include a slot number (member ID). • The dedicated VCPs are **vcp-0** and **vcp-1**. • Ports internal to the VCP subsystem have names corresponding to the PIC and port number. For example, **0/24** indicates internal onboard port 24, and 1/26 indicates internal uplink module port 26.	**brief**
Input Octets/Packets	Total number of octets and packets received on the VCP interface.	**brief member** none
Output Octets/Packets	Total number of octets and packets transmitted on the VCP interface.	**brief member** none
master: ***number***	Member ID of the Virtual Chassis master.	All levels
Port	VCP for which **RX** (Receive) statistics,**TX** (Transmit) statistics, or both are reported by the VCP subsystem during a sampling interval—since the statistics counter was last cleared.	**detail extensive**
Total octets	Total number of octets received and transmitted on the VCP interface.	**detail extensive**
Total packets	Total number of packets received and transmitted on the VCP interface.	**detail extensive**
Unicast packets	Number of unicast packets received and transmitted on the VCP interface.	**detail extensive**
Broadcast packets	Number of broadcast packets received and transmitted on the VCP interface.	**detail extensive**
Multicast packets	Number of multicast packets received and transmitted on the VCP interface.	**detail extensive**
MAC control frames	Number of media access control (MAC) control frames received and transmitted on the VCP interface.	**detail extensive**
CRC alignment errors	Number of packets received on the VCP interface that had a length—excluding framing bits, but including frame check sequence (FCS) octets—of between 64 and 1518 octets, inclusive, and had one of the following errors: • Invalid FCS with an integral number of octets (FCS error) • Invalid FCS with a nonintegral number of octets (alignment error)	**detail extensive**
Oversize packets	Number of packets received on the VCP interface that were longer than 1518 octets (excluding framing bits, but including FCS octets) but were otherwise well formed.	**detail extensive**
Undersize packets	Number of packets received on the VCP interface that were shorter than 64 octets (excluding framing bits but including FCS octets) and were otherwise well formed..	**detail extensive**

Table 116: show virtual-chassis vc-port statistics Output Fields *(continued)*

Field Name	Field Description	Level of Output
Jabber packets	Number of packets received on the VCP interface that were longer than 1518 octets—excluding framing bits, but including FCS octets—and that had either an FCS error or an alignment error. NOTE: This definition of *jabber* is different from the definition in IEEE-802.3 section 8.2.1.5 (10Base5) and section 10.3.1.4 (10Base2). These documents define *jabber* as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is between 20 ms and 150 ms.	**detail extensive**
Fragments received	Number of packets received on the VCP interface that were shorter than 64 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted.	**detail extensive**
Ifout errors	Number of outbound packets received on the VCP interface that could not be transmitted because of errors.	**detail extensive**
Packet drop events	Number of outbound packets received on the VCP interface that were dropped, rather than being encapsulated and sent out of the switch as fragments. The packet drop counter is incremented if a temporary shortage of packet memory causes packet fragmentation to fail.	**detail extensive**
64 octets frames	Number of packets received on the VCP interface (including invalid packets) that were 64 octets in length (excluding framing bits, but including FCS octets).	**detail extensive**
65–127 octets frames	Number of packets received on the VCP interface (including invalid packets) that were between 65 and 127 octets in length, inclusive (excluding framing bits, but including FCS octets).	**detail extensive**
128–255 octets frames	Number of packets received on the VCP interface (including invalid packets) that were between 128 and 255 octets in length, inclusive (excluding framing bits, but including FCS octets).	**detail extensive**
256–511 octets frames	Number of packets received on the VCP interface (including invalid packets) that were between 256 and 511 octets in length, inclusive (excluding framing bits, but including FCS octets).	**detail extensive**
512–1023 octets frames	Number of packets received on the VCP interface (including invalid packets) that were between 512 and 1023 octets in length, inclusive (excluding framing bits, but including FCS octets).	**detail extensive**
1024–1518 octets frames	Number of packets received on the VCP interface (including invalid packets) that were between 1024 and 1518 octets in length, inclusive (excluding framing bits, but including FCS octets).	**detail extensive**
Rate packets per second	Number of packets per second received and transmitted on the VCP interface.	**detail extensive**
Rate bytes per second	Number of bytes per second received and transmitted on the VCP interface.	**detail extensive**

show virtual-chassis vc-port statistics

```
user@SWA-0> show virtual-chassis vc-port statistics
fpc0:
--------------------------------------------------------------------------
Interface             Input Octets/Packets          Output Octets/Packets
internal-0/24         0           / 0               0           / 0
internal-0/25         0           / 0               0           / 0
internal-1/26         0           / 0               0           / 0
internal-1/27         0           / 0               0           / 0
vcp-0                 0           / 0               0           / 0
vcp-1                 0           / 0               0           / 0
internal-0/26         0           / 0               0           / 0
internal-0/27         0           / 0               0           / 0
internal-1/24         0           / 0               0           / 0
internal-1/25         0           / 0               0           / 0

{master:0}
```

show virtual-chassis vc-port statistics brief

```
user@SWA-0> show virtual-chassis vc-port statistics brief
fpc0:
--------------------------------------------------------------------------
Interface             Input Octets/Packets          Output Octets/Packets
internal-0/24         0           / 0               0           / 0
internal-0/25         0           / 0               0           / 0
internal-1/26         0           / 0               0           / 0
internal-1/27         0           / 0               0           / 0
vcp-0                 0           / 0               0           / 0
vcp-1                 0           / 0               0           / 0
internal-0/26         0           / 0               0           / 0
internal-0/27         0           / 0               0           / 0
internal-1/24         0           / 0               0           / 0
internal-1/25         0           / 0               0           / 0

{master:0}
```

show virtual-chassis vc-port statistics extensive

```
user@SWA-0> show virtual-chassis vc-port statistics extensive
fpc0:
--------------------------------------------------------------------------
                              RX                    TX

Port: internal-0/24
Total octets:                 0                     0
Total packets:                0                     0
Unicast packets:              0                     0
Broadcast packets:            0                     0
Multicast packets:            0                     0
MAC control frames:           0                     0
CRC alignment errors:         0
Oversize packets:             0
Undersize packets:            0
Jabber packets:               0
Fragments received:           0
Ifout errors:                 0
Packet drop events:           0
64        octets frames:      0
65-127    octets frames:      0
128-255   octets frames:      0
256-511   octets frames:      0
512-1023  octets frames:      0
1024-1518 octets frames:      0
```

```
Rate packets per second:   0                    0
Rate bytes per second:     0                    0

...

Port: vcp-0
Total octets:              0                    0
Total packets:             0                    0
Unicast packets:           0                    0
Broadcast packets:         0                    0
Multicast packets:         0                    0
MAC control frames:        0                    0
CRC alignment errors:      0
Oversize packets:          0
Undersize packets:         0
Jabber packets:            0
Fragments received:        0
Ifout errors:              0
Packet drop events:        0
64        octets frames:   0
65-127    octets frames:   0
128-255   octets frames:   0
256-511   octets frames:   0
512-1023  octets frames:   0
1024-1518 octets frames:   0
Rate packets per second:   0                    0
Rate bytes per second:     0                    0

Port: vcp-1
Total octets:              0                    0
Total packets:             0                    0
Unicast packets:           0                    0
Broadcast packets:         0                    0
Multicast packets:         0                    0
MAC control frames:        0                    0
CRC alignment errors:      0
Oversize packets:          0
Undersize packets:         0
Jabber packets:            0
Fragments received:        0
Ifout errors:              0
Packet drop events:        0
64        octets frames:   0
65-127    octets frames:   0
128-255   octets frames:   0
256-511   octets frames:   0
512-1023  octets frames:   0
1024-1518 octets frames:   0
Rate packets per second:   0                    0
Rate bytes per second:     0                    0

...

{master:0}
```

show virtual-chassis vc-port statistics member 0

```
user@SWA-0>show virtual-chassis vc-port statistics member 0
fpc0:
--------------------------------------------------------------------------
Interface             Input Octets/Packets          Output Octets/Packets
internal-0/24         0           / 0               0           / 0
internal-0/25         0           / 0               0           / 0
```

```
internal-1/26          0          / 0              0          / 0
internal-1/27          0          / 0              0          / 0
vcp-0                  0          / 0              0          / 0
vcp-1                  0          / 0              0          / 0
internal-0/26          0          / 0              0          / 0
internal-0/27          0          / 0              0          / 0
internal-1/24          0          / 0              0          / 0
internal-1/25          0          / 0              0          / 0

{master:0}
```

PART 12

Interfaces on J-EX Series Switches

CHAPTER 50

Interfaces—Overview

J-EX Series Switches Interfaces Overview

J-EX Series Switches have two types of interfaces: network interfaces and special interfaces. This topic provides brief information on these interfaces. For additional information, see the *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.

For information on interface-naming conventions on J-EX Series Switches, see "Understanding Interface Naming Conventions on J-EX Series Switches" on page 865.

This topic describes:

Network Interfaces

Network interfaces connect to the network and carry network traffic. Table 117 on page 863 lists the types of network interfaces supported on J-EX Series switches.

Table 117: Network Interface Types and Purposes

Type	Purpose
Aggregated Ethernet interfaces	All J-EX Series switches allow you to group Ethernet interfaces at the physical layer to form a single link layer interface, also known as a *link aggregation group (LAG)* or *bundle*. These aggregated Ethernet interfaces help to balance traffic and increase the uplink bandwidth.

Table 117: Network Interface Types and Purposes *(continued)*

Type	Purpose
LAN access interfaces	Use these J-EX Series switch interfaces to connect a personal computer, laptop, file server, or printer to the network. When you power on a J-EX Series switch and use the factory-default configuration, the software automatically configures interfaces in access mode for each of the network ports. The default configuration also enables autonegotiation for both speed and link mode.
Power over Ethernet (PoE) interfaces	J-EX Series switches provide PoE network ports with various switch models. These ports can be used to connect voice over IP (VoIP) telephones, wireless access points, video cameras, and point-of-sale devices to safely receive power from the same access ports that are used to connect personal computers to the network. PoE interfaces are enabled by default in the factory configuration.
Trunk interfaces	J-EX Series access switches can be connected to a distribution switch or customer-edge (CE) switches or routers. To use a port for this type of connection, you must explicitly configure the port interface for trunk mode. The interfaces from the distribution switch or CE switch to the access switches must also be configured for trunk mode.

Special Interfaces

Table 118 on page 864 lists the types of special interfaces supported on J-EX Series switches.

Table 118: Special Interface Types and Purposes

Type	Purpose
Console port	Each J-EX Series switch has a serial port, labeled **CON** or **CONSOLE**, for connecting tty-type terminals to the switch using standard PC-type tty cables. The console port does not have a physical address or IP address associated with it. However, it is an interface in the sense that it provides access to the switch. On J-EX4200 switches that are configured as a Virtual Chassis, you can access the master and configure all members of the Virtual Chassis through any member's console port. For more information on the console port in a Virtual Chassis, see “Understanding Global Management of a Virtual Chassis Configuration” on page 699.
Loopback	All J-EX Series switches have this software-only virtual interface that is always up. The loopback interface provides a stable and consistent interface and IP address on the switch.
Management interface	The Junos OS for J-EX Series switches automatically creates the switch's management Ethernet interface, **me0**. The management Ethernet interface provides an out-of-band method for connecting to the switch. To use **me0** as a management port, you must configure its logical port, **me0.0**, with a valid IP address. You can connect to the management interface over the network using utilities such as SSH or Telnet. SNMP can use the management interface to gather statistics from the switch. (The management interface **me0** is analogous to the **fxp0** interfaces on routers running the Junos OS.)
Routed VLAN Interface (RVI)	J-EX Series switches use a Layer 3 routed VLAN interface (RVI) named **vlan** to route traffic from one broadcast domain to another and to perform other Layer 3 functions such as traffic engineering. These functions are typically performed by a router interface in a traditional network. The RVI functions as a logical router, eliminating the need for having both a switch and a router. The RVI (the **vlan** interface) must be configured as part of a broadcast domain or virtual private LAN service (VPLS) routing instance for Layer 3 traffic to be routed out of it.

Table 118: Special Interface Types and Purposes *(continued)*

Type	Purpose
Virtual Chassis port (VCP) interfaces	Each J-EX4200 switch has two dedicated *Virtual Chassis ports (VCPs)* on its rear panel. These ports can be used to interconnect two to ten J-EX4200 switches as a *Virtual Chassis*, which functions as a single network entity. See "Understanding the High-Speed Interconnection of the Virtual Chassis Members" on page 702. When you power on J-EX Series switches that are interconnected in this manner, the software automatically configures the VCP interfaces for the dedicated ports that have been interconnected. These VCP interfaces are not configurable or modifiable. You can also interconnect J-EX4200 switches across distances of up to 25 miles (40 km) by using the SFP or SFP+ uplink module ports. To do so, you must explicitly set the uplink module ports on the members you want to connect as VCPs. See "Setting an Uplink Module Port as a Virtual Chassis Port (CLI Procedure)" on page 792. When you set the uplink module ports as uplink VCPs and connect member switches through those uplink VCPs, a LAG is automatically formed when the link speed is the same on connected VCPs and at least two VCPs on one member are connected to at least two VCPs on another member. See "Understanding Virtual Chassis Configurations and Link Aggregation" on page 702.
Virtual management Ethernet (VME) interface	J-EX4200 switches have a VME interface. This is a logical interface that is used for Virtual Chassis configurations and allows you to manage all the members of the Virtual Chassis through the master. For more information on the VME interface, see "Understanding Global Management of a Virtual Chassis Configuration" on page 699.

Related Documentation

- J-EX4200 Switches Hardware Overview on page 25
- J-EX8208 Switch Hardware Overview on page 27
- J-EX8216 Switch Hardware Overview on page 30
- PoE and J-EX Series Switches Overview on page 3009
- Understanding Aggregated Ethernet Interfaces and LACP on page 867
- Understanding Layer 3 Subinterfaces on page 871

Understanding Interface Naming Conventions on J-EX Series Switches

J-EX Series Switches use a naming convention for defining the interfaces that is similar to that of other platforms running under the Junos OS. This topic provides brief information on the naming conventions used for interfaces on J-EX Series switches. For additional information, see the *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.

This topic describes:

- Physical Part of an Interface Name on page 865
- Logical Part of an Interface Name on page 866
- Wildcard Characters in Interface Names on page 867

Physical Part of an Interface Name

Interfaces in Junos OS are specified as follows:

type-fpc / pic / port

J-EX Series switches apply this convention as follows:

- *type*—J-EX Series interfaces use the following media types:
  - **ge**—Gigabit Ethernet interface
  - **xe**—10 Gigabit Ethernet interface
  - **fe**—Fast Ethernet interface
- *fpc*—Flexible PIC Concentrator. J-EX Series interfaces use the following convention for the FPC number in interface names:
  - On J-EX4200 standalone switches, FPC refers to the switch itself. The FPC number is always **0** on these switches.
  - On J-EX4200 switches configured in a Virtual Chassis, the FPC number indicates the member ID of the switch within the Virtual Chassis, from **0** through **9**.
  - On J-EX8200 switches, the FPC number indicates the slot number of the line card that contains the physical interface.
- *pic*—J-EX Series interfaces use the following convention for the PIC (Physical Interface Card) number in interface names:
  - On J-EX4200 switches, the PIC number is **0** for all built-in interfaces (interfaces that are not an uplink port) and **1** for uplink ports.
  - On J-EX8200 switches, the PIC number is always **0**.
- *port*—J-EX Series interfaces use the following convention for port numbers:
  - On J-EX4200 switches, built-in network ports are numbered from left to right. On models that have two rows of ports, the ports on the top row start with **0** followed by the remaining even-numbered ports, and the ports on the bottom row start with **1** followed by the remaining odd-numbered ports.
  - Uplink ports in J-EX4200 switches are labeled from left to right, starting with **0**.
  - On J-EX8200 switches, the network ports are numbered from left to right on each line card. On line cards that have two rows of ports, the ports on the top row start with **0** followed by the remaining even-numbered ports, and the ports on the bottom row start with **1** followed by the remaining odd-numbered ports.

Logical Part of an Interface Name

The logical unit part of the interface name corresponds to the logical unit number, which can be a number from 0 through 16384. In the virtual part of the name, a period (.) separates the port and logical unit numbers: *type-fpc/pic/port.logical-unit-number*. For example, if you issue the **show ethernet-switching interfaces** command on a system with a default VLAN, the resulting display shows the logical interfaces associated with the VLAN:

```
Interface   State    VLAN members          Blocking
ge-0/0/0.0  down     remote-analyzer       unblocked
```

```
ge-0/0/1.0  down     default              unblocked
ge-0/0/10.0 down     default              unblocked
```

When you configure aggregated Ethernet interfaces, you configure a logical interface that is called a *bundle* or a *LAG*. Each LAG can include up to 8 or 12 Ethernet interfaces, depending on the switch model.

Wildcard Characters in Interface Names

In the **show interfaces** and **clear interfaces** commands, you can use wildcard characters in the ***interface-name*** option to specify groups of interface names without having to type each name individually. You must enclose all wildcard characters except the asterisk (*) in quotation marks (" ").

Related Documentation

- J-EX Series Switches Interfaces Overview on page 863
- Front Panel of a J-EX4200 Switch
- Slot Numbering for a J-EX8208 Switch
- Slot Numbering for a J-EX8216 Switch

Understanding Aggregated Ethernet Interfaces and LACP

IEEE 802.3ad link aggregation enables you to group Ethernet interfaces to form a single link layer interface, also known as a *link aggregation group (LAG)* or *bundle*.

Aggregating multiple links between physical interfaces creates a single logical point-to-point trunk link or a LAG. The LAG balances traffic across the member links within an aggregated Ethernet bundle and effectively increases the uplink bandwidth. Another advantage of link aggregation is increased availability, because the LAG is composed of multiple member links. If one member link fails, the LAG continues to carry traffic over the remaining links.

Link Aggregation Control Protocol (LACP), a component of IEEE 802.3ad, provides additional functionality for LAGs.

This topic describes:

- Link Aggregation Group (LAG) on page 867
- Link Aggregation Control Protocol (LACP) on page 868

Link Aggregation Group (LAG)

You configure a LAG by specifying the link number as a physical device and then associating a set of interfaces (ports) with the link. All the interfaces must have the same speed and be in full-duplex mode. The Junos OS for J-EX Series Switches assigns a unique ID and port priority to each interface. The ID and priority are not configurable.

The number of interfaces that can be grouped into a LAG and the total number of LAGs supported on a switch varies according to switch model. Table 119 on page 868 lists the J-EX Series switches and the maximum number of interfaces per LAG and maximum number of LAGs they support.

Table 119: Maximum Interfaces per LAG and Maximum LAGs per Switch

Switch Model	Maximum Interfaces per LAG	Maximum LAGs
J-EX4200	8	64
J-EX8200	12	255

When configuring LAGs, consider the following guidelines:

- The LAG must be configured on both sides of the link.
- The interfaces on either side of the link must be set to the same speed.
- You can configure and apply firewall filters on a LAG.
- LACP can optionally be configured for link negotiation.

You can combine physical Ethernet ports belonging to different member switches of a Virtual Chassis configuration to form a LAG. See "Understanding Virtual Chassis Configurations and Link Aggregation" on page 702.

NOTE: The interfaces that are included within a bundle or LAG are sometimes referred to as *member interfaces*. Do not confuse this term with *member switches*, which refers to J-EX4200 Ethernet Switches that are interconnected as a Virtual Chassis. It is possible to create a LAG that is composed of member interfaces that are located in different member switches of a Virtual Chassis.

A LAG creates a single logical point-to-point connection. A typical deployment for a LAG would be to aggregate trunk links between an access switch and a distribution switch or customer edge (CE) router.

Link Aggregation Control Protocol (LACP)

When LACP is configured, it detects misconfigurations on the local end or the remote end of the link.

About enabling LACP:

- When LACP is not enabled, a local LAG might attempt to transmit packets to a remote single interface, which causes the communication to fail.
- When LACP is enabled, a local LAG cannot transmit packets unless a LAG with LACP is also configured on the remote end of the link.

By default, Ethernet links do not exchange protocol data units (PDUs), which contain information about the state of the link. You can configure Ethernet links to actively transmit PDUs, or you can configure the links to passively transmit them, sending out LACP PDUs only when they receive them from another link. The transmitting link is known as the *actor* and the receiving link is known as the *partner*.

In a scenario where a dual-homed server is deployed with a switch, the network interface cards form a LAG with the switch. During a server upgrade, the server may not be able to exchange LACP PDUs. In such a situation you can configure an interface to be in the UP state even if no PDUs are exchanged. Use the **force-up** statement to configure an interface when the peer has limited LACP capability. The interface selects the associated LAG by default, whether the switch and peer are both in active or passive mode. When there are no received PDUs, the partner is considered to be working in the passive mode. Therefore, LACP PDU transmissions are controlled by the transmitting link.

If the remote end of the LAG link is a security device, LACP might not be supported because security devices require a deterministic configuration. In this case, do not configure LACP. All links in the LAG are permanently operational unless the switch detects a link failure within the Ethernet physical layer or data link layers.

Related Documentation

- Understanding Virtual Chassis Configurations and Link Aggregation on page 702
- Understanding Redundant Trunk Links on J-EX Series Switches on page 1049
- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 740
- Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 746
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

Understanding Interface Ranges on J-EX Series Switches

You can use the interface ranges to group interfaces of the same type that share a common configuration profile. This helps reduce the time and effort in configuring interfaces on J-EX Series switches. The configurations common to all the interfaces can be included in the interface range definition.

The interface range definition contains the name of the interface range defined, the names of the individual member interfaces that do not fall in a series of interfaces, a range of interfaces defined in the member range, and the configuration statements common to all the interfaces. An interface range defined with member ranges and individual members but without any common configurations, is also a valid definition.

NOTE: The interface range definition is supported only for Gigabit, 10-Gigabit, and Fast Ethernet interfaces.

The common configurations defined in the interface range will be overridden by the local configuration.

The defined interface ranges can be used at places where the **interface** node is used in the following configuration hierarchies:

- ethernet-switching-options analyzer *name* input egress interface
- ethernet-switching-options analyzer *name* input ingress interface
- ethernet-switching-options analyzer output interface
- ethernet-switching-options bpdu-block interface
- ethernet-switching-options interfaces
- ethernet-switching-options redundant-trunk-group group-name interface
- ethernet-switching-options secure-access-port interface
- ethernet-switching-options voip interface
- poe interface
- protocols dot1x authentication interface
- protocols gvrp interface
- protocols igmp interface
- protocols igmp-snooping vlan *vlan-name* interface
- protocols isis interface
- protocols link-management peer lmp-control-channel interface
- protocols link-management te-link *name* interface
- protocols lldp interface
- protocols lldp-med interface
- protocols mpls interface
- protocols mstp interface
- protocols mstp msti-*id* interface
- protocols mstp msti-*id* vlan *vlan-id* interface
- protocols oam ethernet link-fault-management interface
- protocols ospf area
- protocols pim interface
- protocols rip group *group-name* neighbor
- protocols ripng group *group-name* neighbor
- protocols router-advertisement interface
- protocols router-discovery interface
- protocols rsvp interface
- protocols sflow interfaces
- protocols stp interface

- **protocols vstp vlan *vlan-id* interface**
- **vlans *vlan-name* interface**

Related Documentation

- J-EX Series Switches Interfaces Overview on page 863
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 919
- Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 922
- Configuring a Layer 3 Subinterface (CLI Procedure) on page 930
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/
- **interface-range on page 962**

Understanding Layer 3 Subinterfaces

A Layer 3 subinterface is a logical division of a physical interface that operates at the network level and therefore can receive and forward 802.1Q VLAN tags. You can use Layer 3 subinterfaces to route traffic among multiple VLANs along a single trunk line that connects a J-EX Series Switch to a Layer 2 switch. Only one physical connection is required between the switches. This topology is often called a “router on a stick” or a “one-armed router” when the Layer 3 device is a router.

To create Layer 3 subinterfaces on a J-EX Series switch, you enable VLAN tagging, partition the physical interface into logical partitions, and bind the VLAN ID to the logical interface.

You can partition one physical interface into up to 4094 different subinterfaces, one for each VLAN. We recommend that you use the VLAN ID as the subinterface number when you configure the subinterface. The Junos OS reserves VLAN IDs 0 and 4095.

VLAN tagging places the VLAN ID in the frame header, allowing each physical interface to handle multiple VLANs. When you configure multiple VLANs on an interface, you must also enable tagging on that interface. Junos OS on J-EX Series switches supports a subset of the 802.1Q standard for receiving and forwarding routed or bridged Ethernet frames with single VLAN tags and running Virtual Router Redundancy Protocol (VRRP) over 802.1Q-tagged interfaces. Double-tagging is not supported.

Related Documentation

- J-EX Series Switches Interfaces Overview on page 863
- Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch on page 893
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

Understanding Unicast RPF for J-EX Series Switches

Unicast reverse-path forwarding (RPF) helps protect the switch against denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks by verifying the unicast source address of each packet that arrives on an ingress interface where unicast RPF is enabled. It also helps ensure that traffic arriving on ingress interfaces comes from a network source that the receiving interface can reach.

When you enable unicast RPF, the switch forwards a packet only if the receiving interface is the best return path to the packet's unicast source address. This is known as strict mode unicast RPF.

NOTE: On J-EX4200 Ethernet Switches, the switch applies unicast RPF *globally* to all interfaces when unicast RPF is configured on any interface. For additional information, see "Limitations of the Unicast RPF Implementation on J-EX4200 Switches" on page 875.

This topic covers:

Unicast RPF for J-EX Series Switches Overview

Unicast RPF functions as an ingress filter that reduces the forwarding of IP packets that might be spoofing an address. By default, unicast RPF is disabled on the switch interfaces.

The type of unicast RPF provided on the switches—that is, strict mode unicast RPF is especially useful on untrusted interfaces. An untrusted interface is an interface where untrusted users or processes can place packets on the network segment.

The switch supports only the active paths method of determining the best return path back to a unicast source address. The active paths method looks up the best reverse path entry in the forwarding table. It does not consider alternate routes specified using routing-protocol-specific methods when determining the best return path.

If the forwarding table lists the receiving interface as the interface to use to forward the packet back to its unicast source, it is the best return path interface. Strict mode unicast RPF recognizes only one best return path to a unicast source address.

Use strict mode unicast RPF only on symmetrically routed interfaces. (For information about symmetrically routed interfaces, see "When to Enable Unicast RPF" on page 873.)

For more information about strict unicast RPF, see RFC 3704, *Ingress Filtering for Multihomed Networks* at http://www.ietf.org/rfc/rfc3704.txt.

Unicast RPF Implementation for J-EX Series Switches

This section includes:

Unicast RPF Packet Filtering

When you enable unicast RPF on the switch, the switch handles traffic in the following manner:

- If the switch receives a packet on the interface that is the best return path to the unicast source address of that packet, the switch forwards the packet.
- If the best return path from the switch to the packet's unicast source address is not the receiving interface, the switch discards the packet.
- If the switch receives a packet that has a source IP address that does not have a routing entry in the forwarding table, the switch discards the packet.

Bootstrap Protocol (BOOTP) and DHCP Requests

Bootstrap protocol (BOOTP) and DHCP request packets are sent with a broadcast MAC address and therefore the switch does not perform unicast RPF checks on them. The switch forwards all BOOTP packets and DHCP request packets without performing unicast RPF checks.

Default Route Handling

If the best return path to the source is the default route (**0.0.0.0**) and the default route points to **reject**, the switch discards all unicast RPF packets. If the default route points to a valid network interface, the switch performs a normal unicast RPF check on the packets.

When to Enable Unicast RPF

Enable unicast RPF when you want to ensure that traffic arriving on a network interface comes from a source that resides on a network that that interface can reach. You can enable unicast RPF on untrusted interfaces to filter spoofed packets. For example, a common application for unicast RPF is to help defend an enterprise network from DoS/DDoS attacks coming from the Internet.

Enable unicast RPF only on symmetrically routed interfaces. A symmetrically routed interface uses the same route in both directions between the source and the destination, as shown in Figure 22 on page 874. Symmetrical routing means that if an interface receives a packet, the switch uses the same interface to send a reply to the packet source (the receiving interface matches the forwarding-table entry for the best return path to the source).

Figure 22: Symmetrically Routed Interfaces

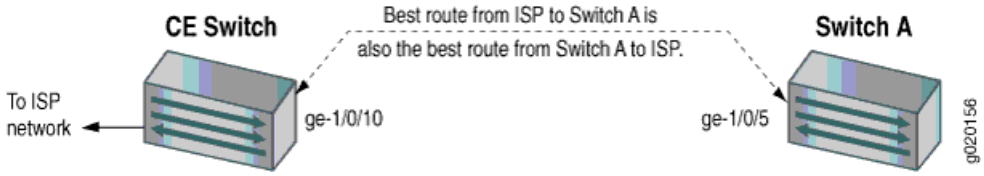

Enabling unicast RPF on asymmetrically routed interfaces (where different interfaces receive a packet and reply to its source) results in packets from legitimate sources being filtered (discarded) because the best return path is not the same interface that received the packet.

The following switch interfaces are most likely to be symmetrically routed and thus are candidates for unicast RPF enabling:

- The service provider edge to a customer
- The customer edge to a service provider
- A single access point out of the network (usually on the network perimeter)
- A terminal network that has only one link

NOTE: Because unicast RPF is enabled globally on J-EX4200 switches, ensure that *all* interfaces are symmetrically routed before you enable unicast RPF on those switches. Enabling unicast RPF on asymmetrically routed interfaces results in packets from legitimate sources being filtered.

TIP: Enabling unicast RPF as close as possible to the traffic source stops spoofed traffic before it can proliferate or reach interfaces that do not have unicast RPF enabled.

When Not to Enable Unicast RPF

Typically, you will not enable unicast RPF if:

- Switch interfaces are multihomed.
- Switch interfaces are trusted interfaces.
- BGP is carrying prefixes and some of those prefixes are not advertised or are not accepted by the ISP under its policy. (The effect in this case is the same as filtering an interface by using an incomplete access list.)
- Switch interfaces face the network core. Core-facing interfaces are usually asymmetrically routed.

An asymmetrically routed interface uses different paths to send and receive packets between the source and the destination, as shown in Figure 23 on page 875. This means that if an interface receives a packet, that interface does not match the forwarding table

entry as the best return path back to the source. If the receiving interface is not the best return path to the source of a packet, unicast RPF causes the switch to discard the packet even though it comes from a valid source.

Figure 23: Asymmetrically Routed Interfaces

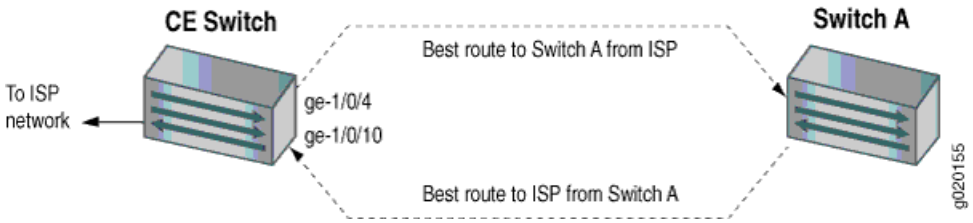

NOTE: Do not enable unicast RPF on J-EX4200 switches if any switch interfaces are asymmetrically routed, because unicast RPF is enabled globally on all interfaces of those switches. All switch interfaces must be symmetrically routed for you to enable unicast RPF without the risk of the switch discarding traffic that you want to forward.

Limitations of the Unicast RPF Implementation on J-EX4200 Switches

On J-EX4200 switches, the switch implements unicast RPF on a global basis. You cannot enable unicast RPF on a per-interface basis. Unicast RPF is globally disabled by default.

- When you enable unicast RPF on any interface, it is automatically enabled on all switch interfaces, including link aggregation groups (LAGs) and routed VLAN interfaces (RVIs).
- When you disable unicast RPF on the interface (or interfaces) on which you enabled unicast RPF, it is automatically disabled on all switch interfaces.

NOTE: You must explicitly disable unicast RPF on every interface on which it was explicitly enabled or unicast RPF remains enabled on all switch interfaces.

The J-EX4200 switches do not perform unicast RPF filtering on equal-cost multipath (ECMP) traffic. The unicast RPF check examines only one best return path to the packet source, but ECMP traffic employs an address block consisting of multiple paths.

Using unicast RPF to filter ECMP traffic on J-EX4200 switches can result in the switch discarding packets that you want to forward because the unicast RPF filter does not examine the entire ECMP address block.

Related Documentation

- Example: Configuring Unicast RPF on a J-EX Series Switch on page 900
- Configuring Unicast RPF (CLI Procedure) on page 927
- Disabling Unicast RPF (CLI Procedure) on page 928

Understanding IP Directed Broadcast for J-EX Series Switches

IP directed broadcast helps you implement remote administration tasks such as backups and wake-on-LAN (WOL) application tasks by sending broadcast packets targeted at the hosts in a specified destination subnet. IP directed broadcast packets traverse the network in the same way as unicast IP packets until they reach the destination subnet. When they reach the destination subnet and IP directed broadcast is enabled on the receiving switch, the switch translates ("explodes") the IP directed broadcast packet into a broadcast that floods the packet on the target subnet. All hosts on the target subnet receive the IP directed broadcast packet.

This topic covers:

IP Directed Broadcast for J-EX Series Switches Overview

IP directed broadcast packets have a destination IP address that is a valid broadcast address for the subnet that is the target of the directed broadcast (the target subnet). The intent of an IP directed broadcast is to flood the target subnet with the broadcast packets without broadcasting to the entire network. IP directed broadcast packets cannot originate from the target subnet.

When you send an IP directed broadcast packet, as it travels to the target subnet, the network forwards it in the same way as it forwards a unicast packet. When the packet reaches a switch that is directly connected to the target subnet, the switch checks to see whether IP directed broadcast is enabled on the interface that is directly connected to the target subnet:

- If IP directed broadcast is enabled on that interface, the switch broadcasts the packet on that subnet by rewriting the destination IP address as the configured broadcast IP address for the subnet. The switch converts the packet to a link-layer broadcast packet that every host on the network processes.
- If IP directed broadcast is disabled on the interface that is directly connected to the target subnet, the switch drops the packet.

IP Directed Broadcast Implementation for J-EX Series Switches

You configure IP directed broadcast on a per-subnet basis by enabling IP directed broadcast on the Layer 3 interface of the subnet's VLAN. When the switch that is connected to that subnet receives a packet that has the subnet's broadcast IP address as the destination address, the switch broadcasts the packet to all hosts on the subnet.

By default, IP directed broadcast is disabled.

When to Enable IP Directed Broadcast

IP directed broadcast is disabled by default. Enable IP directed broadcast when you want to perform remote management or administration services such as backups or WOL tasks on hosts in a subnet that does not have a direct connection to the Internet.

Enabling IP directed broadcast on a subnet affects only the hosts within that subnet. Only packets received on the subnet's Layer 3 interface that have the subnet's broadcast IP address as the destination address are flooded on the subnet.

When Not to Enable IP Directed Broadcast

Typically, you do not enable IP directed broadcast on subnets that have direct connections to the Internet. Disabling IP directed broadcast on a subnet's Layer 3 interface affects only that subnet. If you disable IP directed broadcast on a subnet and a packet that has the broadcast IP address of that subnet arrives at the switch, the switch drops the broadcast packet.

If a subnet has a direct connection to the Internet, enabling IP directed broadcast on it increases the network's susceptibility to denial-of-service (DoS) attacks.

For example, a malicious attacker can spoof a source IP address (use a source IP address that is not the actual source of the transmission to deceive a network into identifying the attacker as a legitimate source) and send IP directed broadcasts containing Internet Control Message Protocol (ICMP) echo (ping) packets. When the hosts on the network with IP directed broadcast enabled receive the ICMP echo packets, they all send replies to the victim that has the spoofed source IP address. This creates a flood of ping replies in a DoS attack that can overwhelm the spoofed source address; this is known as a "smurf" attack. Another common DoS attack on exposed networks with IP directed broadcast enabled is a "fraggle" attack, which is similar to a smurf attack except that the malicious packet is a User Datagram Protocol (UDP) echo packet instead of an ICMP echo packet.

Related Documentation

- Example: Configuring IP Directed Broadcast on a J-EX Series Switch on page 904
- Configuring IP Directed Broadcast (CLI Procedure) on page 929

High Availability Features for J-EX Series Switches Overview

High availability refers to the hardware and software components that provide redundancy and reliability for packet-based communications. This topic covers the following high availability features of J-EX Series Switches:

- VRRP on page 878
- Graceful Protocol Restart on page 878
- Redundant Routing Engines on page 878
- Graceful Routing Engine Switchover on page 879
- Virtual Chassis Software Upgrade and Failover Features on page 879
- Link Aggregation on page 880

VRRP

You can configure the Virtual Router Redundancy Protocol (VRRP) or VRRP for IPv6 on Gigabit Ethernet interfaces, 10-Gigabit Ethernet interfaces, and logical interfaces on J-EX Series switches. When VRRP is configured, the switches act as virtual routing platforms. VRRP enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts. The VRRP routing platforms share the IP address corresponding to the default route configured on the hosts. At any time, one of the VRRP routing platforms is the master (active) and the others are backups. If the master routing platform fails, one of the backup routing platforms becomes the new master, providing a virtual default routing platform and enabling traffic on the LAN to be routed without relying on a single routing platform. Using VRRP, a backup J-EX Series switch can take over a failed default switch within a few seconds. This is done with minimum loss of VRRP traffic and without any interaction with the hosts.

For more information on VRRP, see “Understanding VRRP on J-EX Series Switches” on page 1425.

Graceful Protocol Restart

With standard implementations of routing protocols, any service interruption requires an affected switch to recalculate adjacencies with neighboring switches, restore routing table entries, and update other protocol-specific information. An unprotected restart of a switch can result in forwarding delays, route flapping, wait times stemming from protocol reconvergence, and even dropped packets. Graceful protocol restart allows a restarting switch and its neighbors to continue forwarding packets without disrupting network performance. Because neighboring switches assist in the restart (these neighbors are called helper switches), the restarting switch can quickly resume full operation without recalculating algorithms from scratch.

On J-EX Series switches, graceful protocol restart can be applied to aggregate and static routes and for routing protocols (BGP, IS-IS, OSPF, and RIP).

Graceful protocol restart works similarly for the different routing protocols. The main benefits of graceful protocol restart are uninterrupted packet forwarding and temporary suppression of all routing protocol updates. Graceful protocol restart thus allows a switch to pass through intermediate convergence states that are hidden from the rest of the network. Most graceful restart implementations define two types of switches—the restarting switch and the helper switch. The restarting switch requires rapid restoration of forwarding state information so that it can resume the forwarding of network traffic. The helper switch assists the restarting switch in this process. Individual graceful restart configuration statements typically apply to either the restarting switch or the helper switch.

Redundant Routing Engines

Two to ten J-EX4200 switches can be interconnected to create a Virtual Chassis configuration that operates as a single network entity. Every Virtual Chassis configuration has a master and a backup. The master acts as the master Routing Engine and the backup

acts as the backup Routing Engine. The Routing Engine provides the following functionality:

- Runs various routing protocols
- Provides the forwarding table to the Packet Forwarding Engines (PFEs) in all the member switches of the Virtual Chassis configuration
- Runs other management and control processes for the entire Virtual Chassis configuration

The master Routing Engine, which is in the master of the Virtual Chassis configuration, runs the Junos OS in the master role. It receives and transmits routing information, builds and maintains routing tables, communicates with interfaces and Packet Forwarding Engine components of the member switches, and has full control over the Virtual Chassis configuration.

The backup Routing Engine, which is in the backup of the Virtual Chassis configuration, runs the Junos OS in the backup role. It stays in sync with the master Routing Engine in terms of protocol states, forwarding tables, and so forth. If the master becomes unavailable, the backup Routing Engine takes over the functions that the master Routing Engine performs.

Graceful Routing Engine Switchover

You can configure graceful Routing Engine switchover (GRES) in a Virtual Chassis configuration, allowing the configuration to switch from the master Routing Engine in the master to the backup Routing Engine in the backup with minimal interruption to network communications. When you configure GRES, the backup Routing Engine automatically synchronizes with the master Routing Engine to preserve kernel state information and forwarding state. Any updates to the master Routing Engine are replicated to the backup Routing Engine as soon as they occur. If the kernel on the master Routing Engine stops operating, the master Routing Engine experiences a hardware failure, or the administrator initiates a manual switchover, mastership switches to the backup Routing Engine.

When the backup Routing Engine assumes mastership in a redundant failover configuration (that is, when graceful Routing Engine switchover is not enabled), the Packet Forwarding Engines initialize their state to boot up state before they connect to the new master Routing Engine. In contrast, in a graceful switchover configuration, the Packet Forwarding Engines do not reinitialize their state, but resynchronize their state with the new master Routing Engine. The interruption to the traffic is minimal.

Virtual Chassis Software Upgrade and Failover Features

J-EX4200 switches provide these features for increased resiliency in Virtual Chassis configurations:

- Virtual Chassis atomic software upgrade—When you upgrade software in a Virtual Chassis configuration, the upgrade will either succeed or fail on all member switches, preventing the situation in which only some of the Virtual Chassis member switches are upgraded.

- Virtual Chassis fast failover—A hardware-assisted failover mechanism that automatically reroutes traffic and reduces traffic loss in the event of a link failure.
- Virtual Chassis split and merge—If there is a disruption to the Virtual Chassis configuration due to member switches failing or being removed from the configuration, the Virtual Chassis configuration splits into two separate Virtual Chassis.

Link Aggregation

You can combine multiple physical Ethernet ports to form a logical point-to-point link, known as a *link aggregation group (LAG)* or *bundle*. A LAG provides more bandwidth than a single Ethernet link can provide. Additionally, link aggregation provides network redundancy by load-balancing traffic across all available links. If one of the links should fail, the system automatically load-balances traffic across all remaining links.

You can select up to eight Ethernet interfaces and include them within a LAG. In a J-EX4200 Virtual Chassis configuration, the interfaces that form a LAG can be on different members of the Virtual Chassis. See “Understanding Virtual Chassis Configurations and Link Aggregation” on page 702.

Related Documentation

- For more information on high availability features, see the *Junos OS High Availability Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.
- Virtual Chassis Overview on page 691
- Understanding Virtual Chassis Components on page 694
- Understanding Virtual Chassis Configurations and Link Aggregation on page 702
- Understanding VRRP on J-EX Series Switches on page 1425

CHAPTER 51

Examples: Interfaces Configuration

Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch

J-EX Series switches allow you to combine multiple Ethernet links into one logical interface for higher bandwidth and redundancy. The ports that are combined in this manner are referred to as a link aggregation group (LAG) or bundle. The number of Ethernet links you can combine into a LAG depends on your J-EX Series switch model. See “Understanding Aggregated Ethernet Interfaces and LACP” on page 867 for more information.

This example describes how to configure uplink LAGs to connect a Virtual Chassis access switch to a Virtual Chassis distribution switch:

Requirements

This example uses the following software and hardware components:

- Two J-EX4200-48T switches
- Two J-EX4200-24F switches
- Four uplink modules

Before you configure the LAGs, be sure you have:

- Configured the Virtual Chassis switches. See “Example: Configuring a Virtual Chassis with a Master and Backup in a Single Wiring Closet” on page 717.
- Configured the uplink ports on the switches as trunk ports. See “Configuring Gigabit Ethernet Interfaces (CLI Procedure)” on page 919.

Overview and Topology

For maximum speed and resiliency, you can combine uplinks between an access switch and a distribution switch into LAGs. Using LAGs can be particularly effective when connecting a multimember Virtual Chassis access switch to a multimember Virtual Chassis distribution switch.

The Virtual Chassis access switch in this example is composed of two member switches. Each member switch has an uplink module with two 10-Gigabit Ethernet ports. These ports are configured as trunk ports, connecting the access switch with the distribution switch.

Configuring the uplinks as LAGs has the following advantages:

- Link Aggregation Control Protocol (LACP) can optionally be configured for link negotiation.
- It doubles the speed of each uplink from 10 Gbps to 20 Gbps.
- If one physical port is lost for any reason (a cable is unplugged or a switch port fails, or one member switch is unavailable), the logical port transparently continues to function over the remaining physical port.

The topology used in this example consists of one Virtual Chassis access switch and one Virtual Chassis distribution switch. The access switch is composed of two J-EX4200-48T switches (SWA-0 and SWA-1), interconnected to each other with their Virtual Chassis ports (VCPs) as member switches of Host-A. The distribution switch is composed of two J-EX4200-24F switches (SWD-0 and SWD-1), interconnected with their VCPs as member switches of Host-D.

Each member of the access switch has an uplink module installed. Each uplink module has two ports. The uplinks are configured to act as trunk ports, connecting the access switch with the distribution switch. One uplink port from SWA-0 and one uplink port from SWA-1 are combined as LAG **ae0** to SWD-0. This link is used for one VLAN. The remaining uplink ports from SWA-0 and from SWA-1 are combined as a second LAG connection (**ae1**) to SWD-1. LAG **ae1** is used for another VLAN.

NOTE: If the remote end of the LAG link is a security device, LACP might not be supported because security devices require a deterministic configuration. In this case, do not configure LACP. All links in the LAG are permanently operational unless the switch detects a link failure within the Ethernet physical layer or data link layers.

Figure 24: Topology for LAGs Connecting a Virtual Chassis Access Switch to a Virtual Chassis Distribution Switch

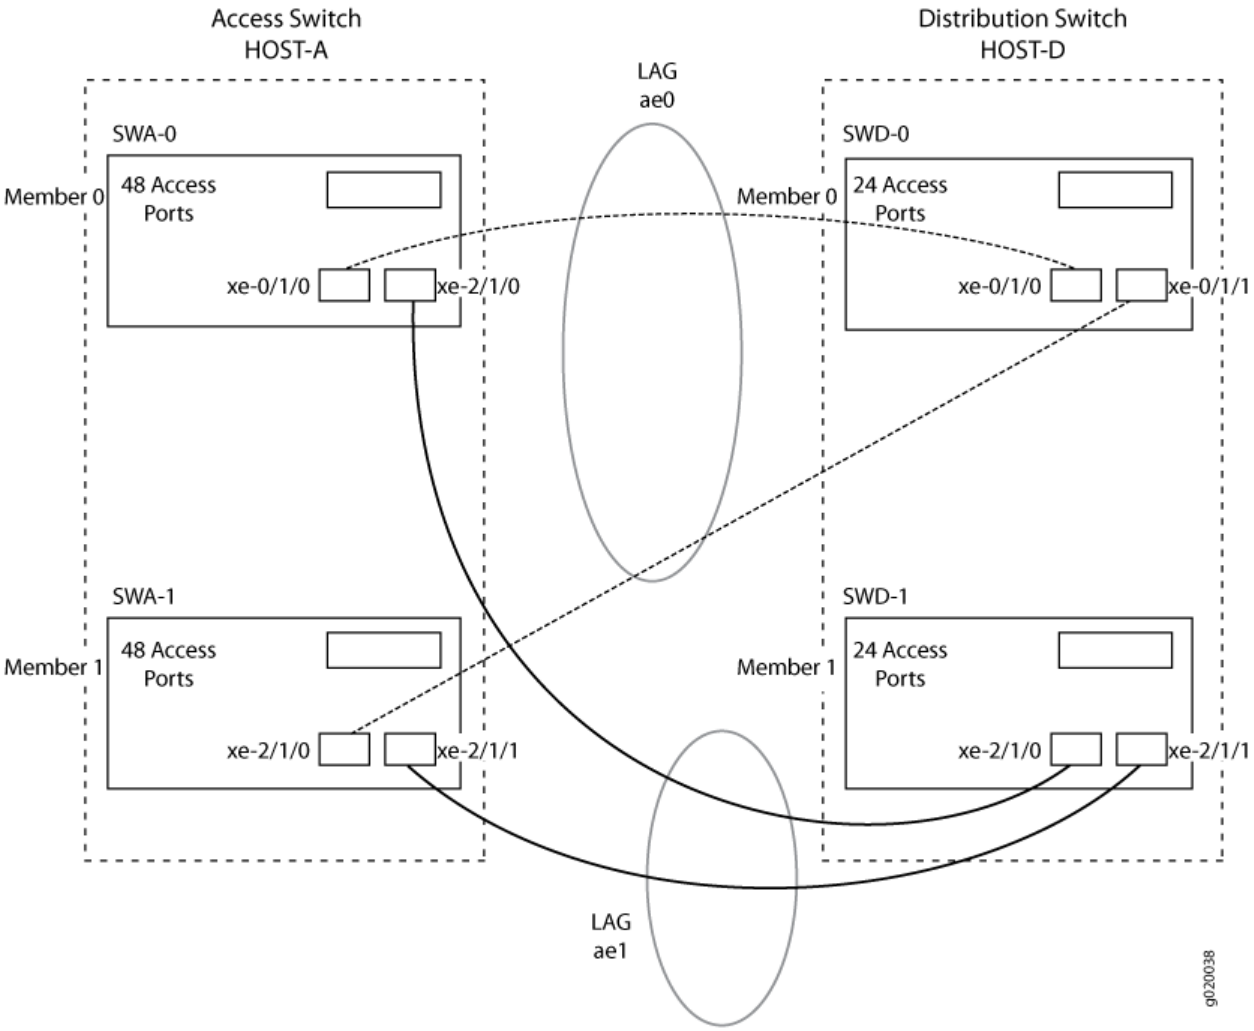

Table 120 on page 883 details the topology used in this configuration example.

Table 120: Components of the Topology for Connecting Virtual Chassis Access Switches to a Virtual Chassis Distribution Switch

Switch	Hostname and VCID	Base Hardware	Uplink Module	Member ID	Trunk Port
SWA-0	Host-A Access switch VCID 1	J-EX4200-48T switch	One uplink module	0	**xe-0/1/0** to SWD-0 **xe-0/1/1** to SWD-1
SWA-1	Host-A Access switch VCID 1	J-EX4200-48P switch	One uplink module	1	**xe-2/1/0** to SWD-0 **xe-2/1/1** to SWD-1

Table 120: Components of the Topology for Connecting Virtual Chassis Access Switches to a Virtual Chassis Distribution Switch *(continued)*

Switch	Hostname and VCID	Base Hardware	Uplink Module	Member ID	Trunk Port
SWD-0	Host-D Distribution switch VCID 4	J-EX4200 L-24F switch	One uplink module	0	**xe-0/1/0** to SWA-0 **xe-0/1/1** to SWA-1
SWD-1	Host-D Distribution switch VCID 4	J-EX4200 L-24F switch	One uplink module	1	**xe-2/1/0** to SWA-0 **xe-2/1/1** to SWA-1

Configuration

To configure two uplink LAGs from the Virtual Chassis access switch to the Virtual Chassis distribution switch:

CLI Quick Configuration

To quickly configure aggregated Ethernet high-speed uplinks between a Virtual Chassis access switch and a Virtual Chassis distribution switch, copy the following commands and paste them into the switch terminal window:

```
[edit]
set chassis aggregated-devices ethernet device-count 2
set interfaces ae0 aggregated-ether-options minimum-links 2
set interfaces ae0 aggregated-ether-options link-speed 10g
set interfaces ae1 aggregated-ether-options minimum-links 2
set interfaces ae1 aggregated-ether-options link-speed 10g
set interfaces ae0 unit 0 family inet address 192.0.2.0/25
set interfaces ae1 unit 1 family inet address 192.0.2.128/25
set interfaces xe-0/1/0 ether-options 802.ad ae0
set interfaces xe-2/1/0 ether-options 802.ad ae0
set interfaces xe-0/1/1 ether-options 802.ad ae1
set interfaces xe-2/1/1 ether-options 802.ad ae1
```

Step-by-Step Procedure

To configure aggregated Ethernet high-speed uplinks between a Virtual Chassis access switch and a Virtual Chassis distribution switch:

1. Specify the number of LAGs to be created on the chassis:

   ```
   [edit chassis]
   user@Host-A# set aggregated-devices ethernet device-count 2
   ```

2. Specify the number of links that need to be present for the **ae0** LAG interface to be **up**:

   ```
   [edit interfaces]
   user@Host-A# set ae0 aggregated-ether-options minimum-links 2
   ```

3. Specify the number of links that need to be present for the **ae1** LAG interface to be **up**:

   ```
   [edit interfaces]
   user@Host-A# set ae1 aggregated-ether-options minimum-links 2
   ```

4. Specify the media speed of the **ae0** link:

```
[edit interfaces]
user@Host-A# set ae0 aggregated-ether-options link-speed 10g
```

5. Specify the media speed of the **ae1** link:

```
[edit interfaces]
user@Host-A# set ae1 aggregated-ether-options link-speed 10g
```

6. Specify the interface ID of the uplinks to be included in LAG **ae0**:

```
[edit interfaces]
user@Host-A# set xe-0/1/0 ether-options 802.ad ae0
user@Host-A# set xe-2/1/0 ether-options 802.ad ae0
```

7. Specify the interface ID of the uplinks to be included in LAG **ae1**:

```
[edit interfaces]
user@Host-A# set xe-0/1/1 ether-options 802.ad ae1
user@Host-A# set xe-2/1/1 ether-options 802.ad ae1
```

8. Specify that LAG **ae0** belongs to the subnet for the employee broadcast domain:

```
[edit interfaces]
user@Host-A# set ae0 unit 0 family inet address 192.0.2.0/25
```

9. Specify that LAG **ae1** belongs to the subnet for the guest broadcast domain:

```
[edit interfaces]
user@Host-A# set ae1 unit 1 family inet address 192.0.2.128/25
```

Results Display the results of the configuration:

```
[edit]
chassis {
    aggregated-devices {
        ethernet {
            device-count 2;
        }
    }
}
interfaces {
    ae0 {
        aggregated-ether-options {
            link-speed 10g;
            minimum-links 2;
        }
        unit 0 {
            family inet {
                address 192.0.2.0/25;
            }
        }
    }
    ae1 {
        aggregated-ether-options {
            link-speed 10g;
            minimum-links 2;
        }
        unit 0 {
            family inet {
```

```
            address 192.0.2.128/25;
          }
        }
      xe–0/1/0 {
        ether-options {
          802.ad ae0;
        }
      }
      xe–2/1/0 {
        ether-options {
          802.ad ae0;
        }
      }
      xe–0/1/1 {
        ether-options {
          802.ad ae1;
        }
      }
      xe–12/1/1 {
        ether-options {
          802.ad ae1;
        }
      }
    }
```

Verification

To verify that switching is operational and two LAGs have been created, perform these tasks:

- Verifying That LAG ae0 Has Been Created on page 886
- Verifying That LAG ae1 Has Been Created on page 886

Verifying That LAG ae0 Has Been Created

Purpose Verify that LAG **ae0** has been created on the switch.

Action **show interfaces ae0 terse**

```
Interface               Admin      Link Proto      Local              Remote

ae0                      up        up
ae0.0                    up        up     inet     10.10.10.2/24
```

Meaning The output confirms that the **ae0** link is up and shows the **family** and IP address assigned to this link.

Verifying That LAG ae1 Has Been Created

Purpose Verify that LAG **ae1** has been created on the switch

Action **show interfaces ae1 terse**

```
Interface               Admin Link Proto    Local                 Remote
ae1                     up    down
ae1.0                   up    down inet
```

Meaning The output shows that the **ae1** link is down.

Troubleshooting

Troubleshooting a LAG That Is Down

Problem The **show interfaces terse** command shows that the LAG is **down:**

Solution Check the following:

- Verify that there is no configuration mismatch.
- Verify that all member ports are up.
- Verify that a LAG is part of family ethernet switching (Layer 2 LAG) or family inet (Layer 3 LAG).
- Verify that the LAG member is connected to the correct LAG at the other end.
- Verify that the LAG members belong to the same switch (or the same Virtual Chassis).

Related Documentation

- Example: Configuring a Virtual Chassis with a Master and Backup in a Single Wiring Closet on page 717
- Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 746
- Example: Connecting an Access Switch to a Distribution Switch on page 1078.
- Virtual Chassis Cabling Configuration Examples for J-EX4200 Switches
- Installing an Uplink Module in a J-EX4200 Switch

Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch

J-EX Series switches allow you to combine multiple Ethernet links into one logical interface for higher bandwidth and redundancy. The ports that are combined in this manner are referred to as a link aggregation group (LAG) or bundle. The number of Ethernet links you can combine into a LAG depends on your J-EX Series switch model. See “Understanding Aggregated Ethernet Interfaces and LACP” on page 867 for more information. J-EX Series switches allow you to further enhance these links by configuring Link Aggregation Control Protocol (LACP).

This example describes how to overlay LACP on the LAG configurations that were created in "Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch" on page 740:

Requirements

This example uses the following software and hardware components:

- Two J-EX4200-48T switches
- Two J-EX4200-24F switches
- Four J-EX Series uplink modules

Before you configure LACP, be sure you have:

- Set up the Virtual Chassis switches. See "Example: Configuring a Virtual Chassis with a Master and Backup in a Single Wiring Closet" on page 717.
- Configured the uplink ports on the switches as trunk ports. See "Configuring Gigabit Ethernet Interfaces (CLI Procedure)" on page 919.
- Configured the LAGs. See "Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch" on page 740

Overview and Topology

This example assumes that you are already familiar with the Example: Configuring Aggregated Ethernet High-Speed Uplinks between Virtual Chassis Access Switch and Virtual Chassis Distribution Switch. The topology in this example is exactly the same as the topology in that other example. This example shows how to use LACP to enhance the LAG functionality.

LACP exchanges are made between *actors* (the transmitting link) and *partners* (the receiving link). The LACP *mode* can be either active or passive.

NOTE: If the actor and partner are both in passive mode, they do not exchange LACP packets, which results in the aggregated Ethernet links not coming up. By default, LACP is in passive mode. To initiate transmission of LACP packets and responses to LACP packets, you must enable LACP in active mode.

By default, the actor and partner send LACP packets every second. You can configure the interval at which the interfaces send LACP packets by including the periodic statement at the **[edit interfaces *interface-name* aggregated-ether-options lacp]** hierarchy level.

The interval can be fast (every second) or slow (every 30 seconds).

Configuring LACP for the LAGs on the Virtual Chassis Access Switch

To configure LACP for the access switch LAGs, perform these tasks:

CLI Quick Configuration

To quickly configure LACP for the access switch LAGs, copy the following commands and paste them into the switch terminal window:

```
[edit]
set interfaces ae0 aggregated-ether-options lacp active periodic fast
set interfaces ae1 aggregated-ether-options lacp active periodic fast
```

Step-by-Step Procedure

To configure LACP for Host-A LAGs **ae0** and **ae1**:

1. Specify the aggregated Ethernet options for both bundles:

```
[edit interfaces]
user@Host-A#set ae0 aggregated-ether-options lacp active periodic fast
user@Host-A#set ae1 aggregated-ether-options lacp active periodic fast
```

Results

Display the results of the configuration:

```
[edit interfaces]
user@Host-A# show
ae0 {
  aggregated-ether-options {
    lacp {
      active;
      periodic fast;
    }
  }
}
ae1 {
  aggregated-ether-options {
    lacp {
      active;
      periodic fast;
    }
  }
}
```

Configuring LACP for the LAGs on the Virtual Chassis Distribution Switch

To configure LACP for the two uplink LAGs from the Virtual Chassis access switch to the Virtual Chassis distribution switch, perform these tasks:

CLI Quick Configuration

To quickly configure LACP for the distribution switch LAGs, copy the following commands and paste them into the switch terminal window:

```
[edit interfaces]
set ae0 aggregated-ether-options lacp passive periodic fast
set ae1 aggregated-ether-options lacp passive periodic fast
```

Step-by-Step Procedure

To configure LACP for Host D LAGs **ae0** and **ae1**:

1. Specify the aggregated Ethernet options for both bundles:

```
[edit interfaces]
user@Host-D#set ae0 aggregated-ether-options lacp passive periodic fast
user@Host-D#set ae1 aggregated-ether-options lacp passive periodic fast
```

Results

Display the results of the configuration:

```
[edit interfaces]
user@Host-D# show
ae0 {
  aggregated-ether-options {
    lacp {
      passive;
      periodic fast;
    }
  }
}
ae1 {
  aggregated-ether-options {
    lacp {
      passive
      periodic fast;
    }
  }
}
```

Verification

To verify that LACP packets are being exchanged, perform these tasks:

- Verifying the LACP Settings on page 890
- Verifying That the LACP Packets Are Being Exchanged on page 891

Verifying the LACP Settings

Purpose

Verify that LACP has been set up correctly.

Action

Use the **show lacp interfaces *interface-name*** command to check that LACP has been enabled as active on one end.

```
user@Host-A> show lacp interfaces xe-0/1/0

Aggregated interface: ae0

    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity

      xe-0/1/0      Actor    No   Yes    No   No   No   Yes     Fast    Active

      xe-0/1/0    Partner    No   Yes    No   No   No   Yes     Fast   Passive

    LACP protocol:   Receive State    Transmit State          Mux State

      xe-0/1/0          Defaulted     Fast periodic           Detached
```

Meaning The output indicates that LACP has been set up correctly and is active at one end.

Verifying That the LACP Packets Are Being Exchanged

Purpose Verify that LACP packets are being exchanged.

Action Use the **show interfaces ae*x* statistics** command to display LACP information.

```
user@Host-A> show interfaces ae0 statistics

Physical interface: ae0, Enabled, Physical link is Down
  Interface index: 153, SNMP ifIndex: 30
  Link-level type: Ethernet, MTU: 1514, Speed: Unspecified, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1,
  Minimum bandwidth needed: 0
  Device flags   : Present Running
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Current address: 02:19:e2:50:45:e0, Hardware address: 02:19:e2:50:45:e0
  Last flapped   : Never
  Statistics last cleared: Never
    Input packets : 0
    Output packets: 0
  Input errors: 0, Output errors: 0

  Logical interface ae0.0 (Index 71) (SNMP ifIndex 34)
    Flags: Hardware-Down Device-Down SNMP-Traps Encapsulation: ENET2
    Statistics        Packets        pps         Bytes          bps
    Bundle:
        Input :             0          0             0            0
        Output:             0          0             0            0
    Protocol inet
      Flags: None
      Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
        Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255
```

Meaning The output here shows that the link is down and that no PDUs are being exchanged.

Troubleshooting

These are some tips for troubleshooting:

Troubleshooting a Nonworking LACP Link

Problem The LACP link is not working.

Solution Check the following:

- Remove the LACP configuration and verify whether the static LAG is up.
- Verify that LACP is configured at both ends.

- Verify that LACP is not passive at both ends.
- Verify whether LACP protocol data units (PDUs) are being exchanged by running the **monitor traffic-interface *lag-member* detail** command.

Related Documentation

- Example: Connecting an Access Switch to a Distribution Switch on page 1078
- Virtual Chassis Cabling Configuration Examples for J-EX4200 Switches
- Installing an Uplink Module in a J-EX4200 Switch

Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch

In a large LAN, you commonly need to partition the network into multiple VLANs. You can configure Layer 3 subinterfaces to route traffic between the VLANs. In one common topology, known as a "router on a stick" or a "one-armed router," you connect a router to an access switch with connections to multiple VLANs.

This example describes how to create Layer 3 subinterfaces on trunk interfaces of a distribution switch and access switch so that you can route traffic among multiple VLANs:

Requirements

This example uses the following hardware and software components:

- For the distribution switch, one J-EX4200-24F switch. This model is designed to be used as a distribution switch for aggregation or collapsed core network topologies and in space-constrained data centers. It has twenty-four 1-Gigabit Ethernet fiber SFP ports and an uplink module with two 10-Gigabit Ethernet ports.
- For the access switch, any Layer 2 switch that supports 802.1Q VLAN tags.

Before you connect the switches, make sure you have:

- Connected the two switches.
- Configured the necessary VLANs. See "Configuring VLANs for J-EX Series Switches (CLI Procedure)" on page 1136 or "Configuring VLANs for J-EX Series Switches (J-Web Procedure)" on page 1133.

Overview and Topology

In a large office with multiple buildings and VLANs, you commonly aggregate traffic from a number of access switches into a distribution switch. This configuration example shows a simple topology to illustrate how to connect a single Layer 2 access switch connected to multiple VLANs to a distribution switch, enabling traffic to pass between those VLANs.

In the example topology, the LAN is segmented into five VLANs, all associated with interfaces on the access switch. One 1-Gigabit Ethernet port on the access switch's uplink module connects to one 1-Gigabit Ethernet port on the distribution switch.

Table 121 on page 894 lists the settings for the example topology.

Table 121: Components of the Topology for Creating Layer 3 Subinterfaces on an Access Switch and a Distribution Switch

Property	Settings
Access switch hardware	Any Layer 2 switch with multiple 1-Gigabit Ethernet ports and at least one 1-Gigabit Ethernet uplink module
Distribution switch hardware	J-EX4200-24F, 24 1-Gigabit Ethernet fiber SPF ports (ge-0/0/0 through ge-0/0/23); one uplink module with 10-Gigabit Ethernet ports
VLAN names and tag IDs	vlan1, tag 101 vlan2, tag 102 vlan3, tag 103 vlan4, tag 104 vlan5, tag 105
VLAN subnets	vlan1: 1.1.1.0/24 (addresses 1.1.1.1 through 1.1.1.254) vlan2: 2.1.1.0/24 (addresses 2.1.1.1 through 2.1.1.254) vlan3: 3.1.1.0/24 (addresses 3.1.1.1 through 3.1.1.254) vlan4: 4.1.1.0/24 (addresses 4.1.1.1 through 4.1.1.254) vlan5: 5.1.1.0/24 (addresses 5.1.1.1 through 5.1.1.254)
Port interfaces	On the access switch: ge-0/1/0 On the distribution switch: ge-0/0/0

Configuring the Access Switch Subinterfaces

CLI Quick Configuration

To quickly create and configure subinterfaces on the access switch, copy the following commands and paste them into the switch terminal window:

```
[edit]
set interfaces ge-0/1/0 vlan-tagging
set interfaces ge-0/1/0 unit 0 vlan-id 101 family inet address 1.1.1.1/24
set interfaces ge-0/1/0 unit 1 vlan-id 102 family inet address 2.1.1.1/24
set interfaces ge-0/1/0 unit 2 vlan-id 103 family inet address 3.1.1.1/24
set interfaces ge-0/1/0 unit 3 vlan-id 104 family inet address 4.1.1.1/24
set interfaces ge-0/1/0 unit 4 vlan-id 105 family inet address 5.1.1.1/24
```

Step-by-Step Procedure

To configure the subinterfaces on the access switch:

1. On the trunk interface of the access switch, enable VLAN tagging:

```
[edit interfaces ge-0/1/0]
user@access-switch# set vlan-tagging
```

2. Bind vlan1's VLAN ID to the logical interface:

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 0 vlan-id 101
```

3. Set vlan1's subinterface IP address:

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 0 family inet address 1.1.1.1/24
```

4. Bind vlan2's VLAN ID to the logical interface:

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 1 vlan-id 102
```

5. Set vlan2's subinterface IP address:

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 1 family inet address 2.1.1.1/24
```

6. Bind vlan3's VLAN ID to the logical interface:

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 2 vlan–id 103
```

7. Set vlan3's subinterface IP address:

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 2 family inet address 3.1.1.1/24
```

8. Bind vlan4's VLAN ID to the logical interface:

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 3 vlan-id 104
```

9. Set vlan4's subinterface IP address:

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 3 family inet address 4.1.1.1/24
```

10. Bind vlan5's VLAN ID to the logical interface:

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 4 vlan-id 105
```

11. Set vlan5's subinterface IP address:

```
[edit interfaces ge-0/1/0]
user@access-switch# set unit 4 family inet address 5.1.1.1/24
```

Results Check the results of the configuration:

```
user@access-switch> show configuration
interfaces {
    ge-0/1/0 {
        vlan-tagging;
        unit 0 {
            vlan-id 101;
            family inet {
                address 1.1.1.1/24;
            }
        }
        unit 1 {
            vlan-id 102;
            family inet {
                address 2.1.1.1/24;
```

```
    }
  }
  unit 2 {
    vlan-id 103;
    family inet {
      address 3.1.1.1/24;
    }
  }
  unit 3 {
    vlan-id 104;
    family inet {
      address 4.1.1.1/24;
    }
  }
  unit 4 {
    vlan-id 105;
    family inet {
      address 5.1.1.1/24;
    }
  }
}
```

Configuring the Distribution Switch Subinterfaces

CLI Quick Configuration

To quickly create and configure subinterfaces on the distribution switch, copy the following commands and paste them into the switch terminal window:

```
[edit]
set interfaces ge-0/0/0 vlan-tagging
set interfaces ge-0/0/0 unit 0 vlan-id 101 family inet address 1.1.1.2/24
set interfaces ge-0/0/0 unit 1 vlan-id 102 family inet address 2.1.1.2/24
set interfaces ge-0/0/0 unit 2 vlan-id 103 family inet address 3.1.1.2/24
set interfaces ge-0/0/0 unit 3 vlan-id 104 family inet address 4.1.1.2/24
set interfaces ge-0/0/0 unit 4 vlan-id 105 family inet address 5.1.1.2/24
```

Step-by-Step Procedure

To configure subinterfaces on the distribution switch:

1. On the trunk interface of the distribution switch, enable VLAN tagging:

   ```
   [edit interfaces ge-0/0/0]
   user@distribution-switch# set vlan-tagging
   ```

2. Bind vlan1's VLAN ID to the logical interface:

   ```
   [edit interfaces ge-0/0/0]
   user@distribution-switch# set unit 0 vlan-id 101
   ```

3. Set vlan1's subinterface IP address:

   ```
   [edit interfaces ge-0/0/0]
   user@distribution-switch# set unit 0 family inet address 1.1.1.2/24
   ```

4. Bind vlan2's VLAN ID to the logical interface:

   ```
   [edit interfaces ge-0/0/0]
   user@distribution-switch# set unit 1 vlan-id 102
   ```

5. Set vlan2's subinterface IP address:

   ```
   [edit interfaces ge-0/0/0]
   ```

```
user@distribution-switch# set unit 1 family inet address 2.1.1.2/24
```

6. Bind vlan3's VLAN ID to the logical interface:

```
[edit interfaces ge-0/0/0]
user@distribution-switch# set unit 2 vlan-id 103
```

7. Set vlan3's subinterface IP address:

```
[edit interfaces ge-0/0/0]
user@distribution-switch# set unit 2 family inet address 3.1.1.2/24
```

8. Bind vlan4's VLAN ID to the logical interface:

```
[edit interfaces ge-0/0/0]
user@distribution-switch# set unit 3 vlan-id 104
```

9. Set vlan4's subinterface IP address:

```
[edit interfaces ge-0/0/0]
user@distribution-switch# set unit 3 family inet address 4.1.1.2/24
```

10. Bind vlan5's VLAN ID to the logical interface:

```
[edit interfaces ge-0/0/0]
user@distribution-switch# set unit 4 vlan-id 105
```

11. Set vlan5's subinterface IP address:

```
[edit interfaces ge-0/0/0]
user@distribution-switch# set unit 4 family inet address 5.1.1.2/24
```

Results

```
user@distribution-switch> show configuration
interfaces {
  ge-0/0/0 {
    vlan-tagging;
    unit 0 {
      vlan-id 101;
      family inet {
        address 1.1.1.2/24;
      }
    }
    unit 1 {
      vlan-id 102;
      family inet {
        address 2.1.1.2/24;
      }
    }
    unit 2 {
      vlan-id 103;
      family inet {
        address 3.1.1.2/24;
      }
    }
    unit 3 {
      vlan-id 104;
      family inet {
        address 4.1.1.2/24;
      }
    }
```

```
    unit 4 {
      vlan-id 105;
      family inet {
        address 5.1.1.2/24;
      }
    }
}
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying That Subinterfaces Were Created on page 898
- Verifying That Traffic Passes Between VLANs on page 898

Verifying That Subinterfaces Were Created

Purpose Verify that the subinterfaces were properly created on the access switch and distribution switch.

Action

1. Use the **show interfaces** command on the access switch:

```
user@access-switch> show interfaces ge-0/1/0 terse

Interface               Admin Link Proto    Local                 Remote
ge-0/1/0                up    up
ge-0/1/0.0              up    up   inet     1.1.1.1/24
ge-0/1/0.1              up    up   inet     2.1.1.1/24
ge-0/1/0.2              up    up   inet     3.1.1.1/24
ge-0/1/0.3              up    up   inet     4.1.1.1/24
ge-0/1/0.4              up    up   inet     5.1.1.1/24
ge-0/1/0.32767          up    up
```

2. Use the show interfaces command on the distribution switch:

```
user@distribution-switch> show interfaces ge-0/0/0 terse

Interface               Admin Link Proto    Local                 Remote
ge-0/0/0                up    up
ge-0/0/0.0              up    up   inet     1.1.1.2/24
ge-0/0/0.1              up    up   inet     2.1.1.2/24
ge-0/0/0.2              up    up   inet     3.1.1.2/24
ge-0/0/0.3              up    up   inet     4.1.1.2/24
ge-0/0/0.4              up    up   inet     5.1.1.2/24
ge-0/0/0.32767          up    up
```

Meaning Each subinterface created is displayed as a **ge-*fpc/pic/port.x*** logical interface, where **x** is the unit number in the configuration. The status is listed as **up**, indicating the link is working.

Verifying That Traffic Passes Between VLANs

Purpose Verify that the distribution switch is correctly routing traffic from one VLAN to another.

Action Ping from the access switch to the distribution switch on each subinterface.

1. From the access switch, ping the address of the vlan1 subinterface on the distribution switch:

```
user@access-switch> ping 1.1.1.2 count 4

PING 1.1.1.2 (1.1.1.2): 56 data bytes
64 bytes from 1.1.1.2: icmp_seq=0 ttl=64 time=0.333 ms
64 bytes from 1.1.1.2: icmp_seq=1 ttl=64 time=0.113 ms
64 bytes from 1.1.1.2: icmp_seq=2 ttl=64 time=0.112 ms
64 bytes from 1.1.1.2: icmp_seq=3 ttl=64 time=0.158 ms

--- 1.1.1.2 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.112/0.179/0.333/0.091 ms
```

2. From the access switch, ping the address of the vlan2 subinterface on the distribution switch:

```
user@access-switch> ping 2.1.1.2 count 4

PING 2.1.1.2 (2.1.1.2): 56 data bytes
64 bytes from 2.1.1.2: icmp_seq=0 ttl=64 time=0.241 ms
64 bytes from 2.1.1.2: icmp_seq=1 ttl=64 time=0.113 ms
64 bytes from 2.1.1.2: icmp_seq=2 ttl=64 time=0.162 ms
64 bytes from 2.1.1.2: icmp_seq=3 ttl=64 time=0.167 ms

--- 2.1.1.2 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.113/0.171/0.241/0.046 ms
```

3. From the access switch, ping the address of the vlan3 subinterface on the distribution switch:

```
user@access-switch> ping 3.1.1.2 count 4

PING 3.1.1.2 (3.1.1.2): 56 data bytes
64 bytes from 3.1.1.2: icmp_seq=0 ttl=64 time=0.341 ms
64 bytes from 3.1.1.2: icmp_seq=1 ttl=64 time=0.162 ms
64 bytes from 3.1.1.2: icmp_seq=2 ttl=64 time=0.112 ms
64 bytes from 3.1.1.2: icmp_seq=3 ttl=64 time=0.208 ms

--- 3.1.1.2 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.112/0.206/0.341/0.085 ms
```

4. From the access switch, ping the address of the vlan4 subinterface on the distribution switch:

```
user@access-switch> ping 4.1.1.2 count 4

PING 4.1.1.2 (4.1.1.2): 56 data bytes
64 bytes from 4.1.1.2: icmp_seq=0 ttl=64 time=0.226 ms
64 bytes from 4.1.1.2: icmp_seq=1 ttl=64 time=0.166 ms
64 bytes from 4.1.1.2: icmp_seq=2 ttl=64 time=0.107 ms
64 bytes from 4.1.1.2: icmp_seq=3 ttl=64 time=0.221 ms

--- 4.1.1.2 ping statistics ---
```

```
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.107/0.180/0.226/0.048 ms
```

5. From the access switch, ping the address of the vlan5 subinterface on the distribution switch:

```
user@access-switch> ping 5.1.1.2 count 4

PING 5.1.1.2 (5.1.1.2): 56 data bytes
64 bytes from 5.1.1.2: icmp_seq=0 ttl=64 time=0.224 ms
64 bytes from 5.1.1.2: icmp_seq=1 ttl=64 time=0.104 ms
64 bytes from 5.1.1.2: icmp_seq=2 ttl=64 time=0.102 ms
64 bytes from 5.1.1.2: icmp_seq=3 ttl=64 time=0.170 ms

--- 5.1.1.2 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.102/0.150/0.224/0.051 ms
```

Meaning If all the ping packets are transmitted and are received by the destination address, the subinterfaces are up and working.

Related Documentation

- Example: Connecting an Access Switch to a Distribution Switch on page 1078
- Configuring a Layer 3 Subinterface (CLI Procedure) on page 930

Example: Configuring Unicast RPF on a J-EX Series Switch

Unicast reverse-path forwarding (RPF) helps protect the switch against denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks by verifying the unicast source address of each packet that arrives on an ingress interface where unicast RPF is enabled.

This example shows how to help defend the switch ingress interfaces against denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks by configuring unicast reverse-path forwarding (RPF) on a customer-edge interface to filter incoming traffic:

Requirements

This example uses the following software and hardware components:

- Two J-EX8200 switches

Before you begin, be sure you have:

- Connected the two switches by symmetrically routed interfaces.
- Ensured that the interface on which you will configure unicast RPF is symmetrically routed.

Overview and Topology

Large amounts of unauthorized traffic such as attempts to flood a network with fake (bogus) service requests in a denial-of-service (DoS) attack can consume network resources and deny service to legitimate users. One way to help prevent DoS and distributed denial-of-service (DDoS) attacks is to verify that incoming traffic originates from legitimate network sources.

Unicast RPF helps ensure that a traffic source is legitimate (authorized) by comparing the source address of each packet that arrives on an interface to the forwarding-table entry for its source address. If the switch uses the same interface that the packet arrived on to reply to the packet's source, this verifies that the packet originated from an authorized source, and the switch forwards the packet. If the switch does not use the same interface that the packet arrived on to reply to the packet's source, the packet might have originated from an unauthorized source, and the switch discards the packet.

This example uses two J-EX8200 switches. On J-EX4200 switches, you cannot configure individual interfaces for unicast RPF. On J-EX4200 switches, the switch applies unicast RPF globally to all interfaces on the switch. See “Understanding Unicast RPF for J-EX Series Switches” on page 872 for more information on limitations regarding the configuration of unicast RPF on J-EX4200 switches.

In this example, an enterprise network's system administrator wants to protect Switch A against potential DoS and DDoS attacks from the Internet. The administrator configures unicast RPF on interface **ge-1/0/10** on Switch A. Packets arriving on interface **ge-1/0/10** on Switch A from the Switch B source also use incoming interface **ge-1/0/10** as the best return path to send packets back to the source.

The topology of this configuration example uses two J-EX8200 switches, Switch A and Switch B, connected by symmetrically routed interfaces:

- Switch A is on the edge of an enterprise network. The interface **ge-1/0/10** on Switch A connects to the interface **ge-1/0/5** on Switch B.
- Switch B is on the edge of the service provider network that connects the enterprise network to the Internet.

Configuration

To enable unicast RPF, perform these tasks:

CLI Quick Configuration

To quickly configure unicast RPF on Switch A, copy the following command and paste it into the switch terminal window:

```
[edit interfaces]
set ge-1/0/10 unit 0 family inet rpf-check
```

Step-by-Step Procedure

To configure unicast RPF on Switch A:

1. Enable unicast RPF on interface **ge-1/0/10**:

```
[edit interfaces]
user@switch# set ge-1/0/10 unit 0 family inet rpf-check
```

Results

Check the results:

```
[edit interfaces]
user@switch# show
ge-1/0/10 {
  unit 0 {
    family inet {
      rpf-check;
    }
  }
}
```

Verification

To confirm that the configuration is correct, perform these tasks:

- Verifying That Unicast RPF Is Enabled on the Switch on page 902

Verifying That Unicast RPF Is Enabled on the Switch

Purpose Verify that unicast RPF is enabled.

Action Verify that unicast RPF is enabled on interface **ge-1/0/10** by using the **show interfaces ge-1/0/10 extensive** or **show interfaces ge-1/0/10 detail** command.

```
user@switch> show interfaces ge-1/0/10 extensive
Physical interface: ge-1/0/10, Enabled, Physical link is Down
  Interface index: 139, SNMP ifIndex: 58, Generation: 140
  Link-level type: Ethernet, MTU: 1514, Speed: Auto, MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled,
  Auto-negotiation: Enabled, Remote fault: Online
  Device flags   : Present Running
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:19:e2:50:95:ab, Hardware address: 00:19:e2:50:95:ab
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   0                   0 bps
   Output bytes  :                   0                   0 bps
   Input  packets:                   0                   0 pps
   Output packets:                   0                   0 pps
   IPv6 transit statistics:
    Input  bytes  :                  0
    Output bytes  :                  0
    Input  packets:                  0
    Output packets:                  0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
```

```
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  FIFO errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets

  0 best-effort                    0                    0                    0

  1 assured-forw                   0                    0                    0

  5 expedited-fo                   0                    0                    0

  7 network-cont                   0                    0                    0

Active alarms  : LINK
Active defects : LINK
MAC statistics:                       Receive         Transmit
  Total octets                              0                0
  Total packets                             0                0
  Unicast packets                           0                0
  Broadcast packets                         0                0
  Multicast packets                         0                0
  CRC/Align errors                          0                0
  FIFO errors                               0                0
  MAC control frames                        0                0
  MAC pause frames                          0                0
  Oversized frames                          0
  Jabber frames                             0
  Fragment frames                           0
  VLAN tagged frames                        0
  Code violations                           0
Filter statistics:
  Input packet count                        0
  Input packet rejects                      0
  Input DA rejects                          0
  Input SA rejects                          0
  Output packet count                                        0
  Output packet pad count                                    0
  Output packet error count                                  0
  CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Incomplete
Packet Forwarding Engine configuration:
  Destination slot: 1

Logical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 59) (Generation 135)
  Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
  Traffic statistics:
   Input  bytes  :                   0
   Output bytes  :                   0
   Input  packets:                   0
   Output packets:                   0
   IPv6 transit statistics:
    Input  bytes  :                  0
    Output bytes  :                  0
    Input  packets:                  0
    Output packets:                  0
  Local statistics:
```

```
  Input  bytes  :                    0
  Output bytes  :                    0
  Input  packets:                    0
  Output packets:                    0
 Transit statistics:
  Input  bytes  :                    0                    0 bps
  Output bytes  :                    0                    0 bps
  Input  packets:                    0                    0 pps
  Output packets:                    0                    0 pps
  IPv6 transit statistics:
   Input  bytes  :                   0
   Output bytes  :                   0
   Input  packets:                   0
   Output packets:                   0
     Protocol inet, Generation: 144, Route table: 0
   Flags: uRPF
   Addresses, Flags: Is-Preferred Is-Primary
```

Meaning The second-to-last line of the display shows the unicast RPF flag enabled, confirming that unicast RPF is enabled on interface **ge-1/0/10**.

Related Documentation

- Configuring Unicast RPF (CLI Procedure) on page 927
- Disabling Unicast RPF (CLI Procedure) on page 928

Example: Configuring IP Directed Broadcast on a J-EX Series Switch

IP directed broadcast provides a method of sending broadcast packets to hosts on a specified subnet without broadcasting those packets to hosts on the entire network.

This example shows how to enable a subnet to receive IP directed broadcast packets so you can perform backups and other network management tasks remotely:

Requirements

This example uses the following software and hardware components:

- One PC
- One J-EX Series switch

Before you configure IP directed broadcast for a subnet:

- Ensure that the subnet does not have a direct connection to the Internet.
- Configure routed VLAN interfaces (RVIs) for the ingress and egress VLANs on the switch. See “Configuring Routed VLAN Interfaces (CLI Procedure)” on page 1137 or “Configuring VLANs for J-EX Series Switches (J-Web Procedure)” on page 1133.

Overview and Topology

You might want to perform remote administration tasks such as backups and wake-on-LAN (WOL) application tasks to manage groups of clients on a subnet. One way to do this is to send IP directed broadcast packets targeted at the hosts in a particular target subnet.

The network forwards IP directed broadcast packets as if they were unicast packets. When the IP directed broadcast packet is received by a VLAN that is enabled for **targeted-broadcast**, the switch broadcasts the packet to all the hosts in its subnet.

In this topology (see Figure 25 on page 905), a host is connected to an interface on a J-EX Series switch to manage the clients in subnet **10.1.2.1/24**. When the switch receives a packet with the broadcast IP address of the target subnet as its destination address, it forwards the packet to the subnet's Layer 3 interface and broadcasts it to all the hosts within the subnet.

Figure 25: Topology for IP Directed Broadcast

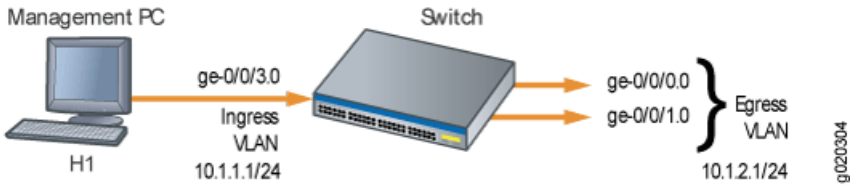

Table 122 on page 905 shows the settings of the components in this example.

Table 122: Components of the IP Directed Broadcast Topology

Property	Settings
Switch hardware	J-EX Series switch
Ingress VLAN name	**v0**
Ingress VLAN IP address	**10.1.1.1/24**
Egress VLAN name	**v1**
Egress VLAN IP address	**10.1.2.1/24**
Interfaces in VLAN **v0**	**ge-0/0/3.0**
Interfaces in VLAN **v1**	**ge-0/0/0.0** and **ge-0/0/1.0**

Configuration

To configure IP directed broadcast on a subnet to enable remote management of its hosts:

CLI Quick Configuration

To quickly configure the switch to accept IP directed broadcasts targeted at subnet **10.1.2.1/24**, copy the following commands and paste them into the switch's terminal window:

```
[edit]
set interfaces ge-0/0/0.0 family ethernet-switching vlan members v1
set interfaces ge-0/0/1.0 family ethernet-switching vlan members v1
set interfaces vlan.1 family inet address 10.1.2.1/24
set interfaces ge-0/0/3.0 family ethernet-switching vlan members v0
set interfaces vlan.0 family inet address 10.1.1.1/24
set vlans v1 l3-interface vlan.1
set vlans v0 l3-interface vlan.0
set interfaces vlan.1 family inet targeted-broadcast
```

Step-by-Step Procedure

To configure the switch to accept IP directed broadcasts targeted at subnet **10.1.2.1/24**:

1. Add logical interface **ge-0/0/0.0** to VLAN **v1**:

   ```
   [edit interfaces]
   user@switch# set ge-0/0/0.0 family ethernet-switching vlan members v1
   ```

2. Add logical interface **ge-0/0/1.0** to VLAN **v1**:

   ```
   [edit interfaces]
   user@switch# set ge-0/0/1.0 family ethernet-switching vlan members v1
   ```

3. Configure the IP address for the egress VLAN, **v1**:

   ```
   [edit interfaces]
   user@switch# set vlan.1 family inet address 10.1.2.1/24
   ```

4. Add logical interface **ge-0/0/3.0** to VLAN **v0**:

   ```
   [edit interfaces]
   user@switch# set ge-0/0/3.0 family ethernet-switching vlan members v0
   ```

5. Configure the IP address for the ingress VLAN:

   ```
   [edit interfaces]
   user@switch# set vlan.0 family inet address 10.1.1.1/24
   ```

6. To route traffic between the ingress and egress VLANs, associate a Layer 3 interface with each VLAN:

   ```
   [edit vlans]
   user@switch# set v1 l3-interfacevlan.1
   user@switch# set v0 l3–interface vlan.0
   ```

7. Enable the Layer 3 interface for the egress VLAN to receive IP directed broadcasts:

   ```
   [edit interfaces]
   user@switch# set vlan.1 family inet targeted-broadcast
   ```

Results

Check the results:

```
user@switch# show
interfaces {
  ge-0/0/0 {
    unit 0 {
      family ethernet-switching {
        vlan {
```

```
            members v1;
          }
        }
      }
    }
    ge-0/0/1 {
      unit 0 {
        family ethernet-switching {
          vlan {
            members v1;
          }
        }
      }
    }
    ge-0/0/3 {
      unit 0 {
        family ethernet-switching {
          vlan {
            members v0;
          }
        }
      }
    }
    vlan {
      unit 0 {
        family inet {
          targeted-broadcast;
          address 10.1.1.1/24;
        }
      }
      unit 1 {
        family inet {
          targeted-broadcast;
          address 10.1.2.1/24;
        }
      }
    }
  vlans {
    default;
    v0 {
      l3-interface vlan.0;
    }
    v1 {
      l3-interface vlan.1;
    }
  }
```

Related Documentation

- Configuring IP Directed Broadcast (CLI Procedure) on page 929

CHAPTER 52

Configuring Interfaces

Configuring Gigabit Ethernet Interfaces (J-Web Procedure)

An Ethernet interface must be configured for optimal performance in a high-traffic network.

To configure properties on a Gigabit Ethernet interface or a 10-Gigabit Ethernet interface on a J-EX Series switch:

1. Select **Interfaces** > **Ports**.

   The page lists Gigabit Ethernet and 10-Gigabit Ethernet interfaces and their link status.

   NOTE: After you make changes to the configuration in this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options** > **Commit**. See "Using the Commit Options to Commit Configuration Changes (J-Web Procedure)" on page 334 for details about all commit options.

2. Select the interface you want to configure. If the interface you want to configure is not listed under **Ports** in the top table on the page, select the FPC (the FPC is the line

card on a J-EX8200 switch or the member switch in a Virtual Chassis configuration) that includes that interface from the **List Ports for FPC** list.

Details for the selected interface such as administrative status, link status, speed, duplex, and flow control are displayed in the bottom table on the page.

NOTE: You can select multiple interfaces and modify their settings at the same time. When you do this, you cannot modify the IP address or enable or disable the administrative status of the selected interface.

3. Click **Edit** and select the set of options you want to configure first:

- Port Role—Enables you to assign a profile for the selected interface.

NOTE: When you select a particular port role, pre-configured port security parameters are set for the VLAN that the interface belongs to. For example, if you select the port role **Desktop**, the port security options **examine-dhcp** and **arp-inspection** are enabled on the VLAN that the interface belongs to. If there are interfaces in the VLAN that have static IP addresses, those interfaces might lose connectivity because those static IP addresses might not be present in the DHCP pool. Therefore, when you are selecting a port role, ensure that the corresponding port security settings for the VLAN are applicable to the interface.

For basic information on port security features such as DHCP snooping (CLI option **examine-dhcp**) or dynamic ARP inspection (DAI) (CLI option **arp-inspection**), see “Configuring Port Security (J-Web Procedure)” on page 2627. For detailed descriptions of port security features, see “Port Security for J-EX Series Switches Overview” on page 2545.

  Click **Details** to view the configuration parameters for the selected port role.

- VLAN Options—Enables you to configure VLAN options for the selected interface.
- Link Options—Enables you to modify the following link options for the selected interface:
  - Speed
  - MTU
  - Autonegotiation
  - Flow Control
  - Duplex
- IP Options—Enables you to configure an IP address for the interface.

4. Configure the interface by configuring options in the selected option set. See Table 123 on page 911 for details on options.
5. Repeat Steps 3 and 4 for the remaining option sets that you want to configure for the interface.

NOTE: To enable or disable the administrative status for a selected interface, click Enable Port or Disable Port.

Table 123: Port Edit Options

Field	Function	Your Action
Port Role	Specifies a profile (role) to assign to the interface. NOTE: Once a port role is configured on the interface, you cannot specify VLAN options or IP options. NOTE: Only the following port roles can be applied on J-EX8200 switch interfaces: • Default • Layer 2 uplink • Routed uplink	
Default	Applies the default role. The interface family is set to **ethernet-switching**, port mode is set to **access**, and RSTP is enabled.	1. Click **Details** to view CLI commands for this role. 2. Click **OK**.
Desktop	Applies the desktop role. The interface family is set to **ethernet-switching**, port mode is set to **access**, RSTP is enabled with the **edge** and **point-to-point** options, and port security parameters (MAC limit =1; dynamic ARP inspection and DHCP snooping enabled) are set.	1. Select an existing VLAN configuration or type the name of a new VLAN configuration to be associated with the interface. 2. Click **Details** to view CLI commands for this role. 3. Click **OK**.
Desktop and Phone	Applies the desktop and phone role. The interface family is set to **ethernet-switching**, port mode is set to **access**, port security parameters (MAC limit =1; dynamic ARP Inspection and DHCP snooping enabled) are set, and recommended CoS parameters are specified for forwarding classes, schedulers, and classifiers. See Table 124 on page 914 for more CoS information.	1. Select an existing VLAN configuration or type the name of a new VLAN configuration to be associated with the interface. You can also select an existing VoIP VLAN configuration or a new VoIP VLAN configuration to be associated with the interface. NOTE: VoIP is not supported on J-EX8200 switches. 2. Click **Details** to view CLI commands for this role. 3. Click **OK**.

Table 123: Port Edit Options *(continued)*

Field	Function	Your Action
Wireless Access Point	Applies the wireless access point role. The interface family is set to **ethernet-switching**, port mode is set to **access**, and RSTP is enabled with the **edge** and **point-to-point** options.	1. Select an existing VLAN configuration or type the name of a new VLAN configuration to be associated with the interface. Type the **VLAN ID** for a new VLAN. 2. Click **Details** to view CLI commands for this role. 3. Click **OK**.
Routed Uplink	Applies the routed uplink role. The interface family is set to **inet**, and recommended CoS parameters are set for schedulers and classifiers. See Table 124 on page 914 for more CoS information.	To specify an IPv4 address: 1. Select the check box **IPv4 address**. 2. Type an IP address—for example: **10.10.10.10**. 3. Enter the subnet mask or address prefix. For example, 24 bits represents **255.255.255.0**. 4. Click **OK**. To specify an IPv6 address: 1. Select the check box **IPv6 address**. 2. Type an IP address—for example: **2001:ab8:85a3::8a2e:370:7334**. 3. Enter the subnet mask or address prefix. 4. Click **OK**.
Layer 2 Uplink	Applies the Layer 2 uplink role. The interface family is set to **ethernet-switching**, port mode is set to **trunk**, RSTP is enabled with the **point-to-point** option, and port security is set to **dhcp-trusted**.	1. For this port role you can select a VLAN member and associate a native VLAN with the interface. 2. Click **Details** to view CLI commands for this role. 3. Click **OK**.
None	Specifies that no port role is configured for the selected interface.	

NOTE: See “Port Role Configuration with the J-Web InterfaceCLI Reference” on page 915 for details on the CLI commands that are associated with each port role.

NOTE: For a J-EX8200 switch, dynamic ARP inspection and DHCP snooping parameters are not configured.

VLAN Options

Table 123: Port Edit Options *(continued)*

Field	Function	Your Action
Port Mode	Specifies the mode of operation for the interface: trunk or access.	If you select **Trunk**, you can: 1. Click **Add** to add a VLAN member. 2. Select the VLAN and click **OK**. 3. (Optional) Associate a native VLAN with the interface. If you select **Access**, you can: 1. Select the VLAN member to be associated with the interface. 2. (Optional) Associate a VoIP VLAN with the interface. Only a VLAN with a VLAN ID can be associated as a VoIP VLAN. NOTE: VoIP is not supported on J-EX8200 switches. Click **OK**.
Link Options		
MTU (bytes)	Specifies the maximum transmission unit size for the interface.	Type a value from **256** through **9216**. The default MTU for Gigabit Ethernet interfaces is **1514**.
Speed	Specifies the speed for the mode.	Select one of the following values: 10 Mbps, 100 Mbps, 1000 Mbps, or Auto-Negotiation.
Duplex	Specifies the link mode.	Select one: **automatic**, **half**, or **full**.
Description	Describes the link. NOTE: If the interface is part of a link aggregation group (LAG), only the option **Description** is enabled.	Enter a brief description for the link.
Enable Auto Negotiation	Enables or disables autonegotiation.	Select the check box to enable autonegotiation, or clear the check box to disable it. By default, autonegotiation is enabled.
Enable Flow Control	Enables or disables flow control.	Select the check box to enable flow control to regulate the amount of traffic sent out of the interface, or clear the check box to disable flow control and permit unrestricted traffic. Flow control is enabled by default.
IP Options		

Table 123: Port Edit Options *(continued)*

Field	Function	Your Action
IPv4 Address	Specifies an IPv4 address for the interface. NOTE: If the IP address is cleared, the interface still belongs to the **inet** family.	1. To specify an IPv4 address, select the check box **IPv4 address**. 2. Type an IP address—for example: **10.10.10.10**. 3. Enter the subnet mask or address prefix. For example, 24 bits represents **255.255.255.0**. 4. Click **OK**.
IPv6 Address	Specifies an IPv6 address for the interface. NOTE: If the IP address is cleared, the interface still belongs to the **inet** family.	1. To specify an IPv6 address, select the check box **IPv6 address**. 2. Type an IP address—for example: **2001:ab8:85a3::8a2e:370:7334**. 3. Enter the subnet mask or address prefix. 4. Click **OK**.

Table 124: Recommended CoS Settings for Port Roles

CoS Parameter	Recommended Settings
Forwarding Classes	There are four forwarding classes: • **voice**—Queue number is set to 7. • **expedited-forwarding**—Queue number is set to 5. • **assured-forwarding**—Queue number is set to 1. • **best-effort**—Queue number is set to 0.
Schedulers	The schedulers and their settings are: • Strict-priority—Transmission rate is set to 10 percent and buffer size to 5 percent. • Expedited-scheduler—Transmission rate is set to 30 percent, buffer size to 30 percent, and priority to **low**. • Assured-scheduler—Transmission rate is set to 25 percent, buffer size to 25 percent, and priority to **low**. • Best-effort scheduler—Transmission rate is set to 35 percent, buffer size to 40 percent, and priority to **low**.
Scheduler maps	When a desktop and phone, routed uplink, or layer 2 uplink role is applied on an interface, the forwarding classes and schedulers are mapped using the scheduler map.
ieee-802.1 classifier	Imports the default **ieee-802.1** classifier configuration and sets the loss priority to **low** for the code point 101 for the **voice** forwarding class.
dscp classifier	Imports the default **dscp** classifier configuration and sets the loss priority to **low** for the code point 101110 for the **voice** forwarding class.

Related Documentation

- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 919
- Monitoring Interface Status and Traffic on page 931

- J-EX Series Switches Interfaces Overview on page 863
- Junos OS CoS for J-EX Series Switches Overview on page 2860
- Understanding Interface Naming Conventions on J-EX Series Switches on page 865

Port Role Configuration with the J-Web Interface (with CLI References)

When you configure Gigabit Ethernet interface properties with the J-Web interface (Configure > Interfaces) you can optionally select pre-configured port roles for those interfaces. When you select a role from the **Port Role** field and apply it to a port, the J-Web interface modifies the switch configuration using CLI commands. Table 125 on page 915 lists the CLI commands applied for each port role.

NOTE: If there is an existing port role configuration, it is cleared before the new port role configuration is applied.

Table 125: Port Role Configuration Summary

Configuration Description	CLI Commands
Default Port Role	
Set the port role to **Default**.	`set interfaces interfaceapply-macro juniper-port-profile Default`
Set port family to **ethernet-switching**. Set port mode to **access**.	`set interfaces interface unit 0 family ethernet-switching port-mode access`
Enable RSTP if redundant trunk groups are not configured.	`delete protocols rstp interface interface disable`
Disable RSTP if redundant trunk groups are configured.	`set protocols rstp interface interface disable`
Desktop Port Role	
Set the port role to desktop.	`set interfaces interface apply-macro juniper-port-profile Desktop`
Set VLAN if new VLAN is specified.	set vlans <*vlan name*> vlan-id <*vlan-id*>
Set port family to **ethernet-switching**. Set Port Mode to **Access**.	`set interfaces interface unit 0 family ethernet-switching port-mode access`
Set VLAN if new VLAN is specified.	`set interfaces interface unit 0 family ethernet-switching vlan members vlan-members`
Set port security parameters.	`set ethernet-switching-options secure-access-port vlan MacTest arp-inspection`
Set RSTP protocol with **edge** option.	`set protocols rstp interface interface edge`

Table 125: Port Role Configuration Summary *(continued)*

Configuration Description	CLI Commands
RSTP protocol is disabled if redundant trunk groups are configured.	`set protocols rstp interface interface disable`
Desktop and Phone Port Role	
Set the port role to desktop and phone.	`set interfaces interfaceapply-macro juniper-port-profile Desktop and Phone`
Set data VLAN if new VLAN is specified. Set voice VLAN if new voice VLAN is specified.	`set vlans vlan-namevlan-id vlan id`
Set port family to **ethernet-switching.** Set Port Mode to **access**.	`set interfaces interfaceunit 0 family ethernet-switching port-mode access`
Set data VLAN on port stanza.	`set interfaces interface unit 0 family ethernet-switching vlan members vlan-members`
Set port security parameters.	`set ethernet-switching-options secure-access-port vlan MacTest arp-inspection`
Set VOIP VLAN.	`set ethernet-switching-options voip interface interface.0 vlan vlan vlan name`

Set class of service parameters
 	
SCHEDULER_MAP=juniper-port-profile-map
 	
IEEE_CLASSIFIER=juniper-ieee-classifier
 	
DSCP_CLASSIFIER=juniper-dscp-classifier
	`set class-of-service interfaces interfacescheduler-map juniper-port-profile-map`
	`set class-of-service interfaces interface unit 0 classifiers ieee-802.1 juniper_ieee_classifier`
	`set class-of-service interfaces interfaceunit 0 classifiers dscp juniper-dscp-classifier`

Set CoS Configuration	Refer to Table 126 on page 918 for details.
Wireless Access Point Port Role	
Set the port role to wireless access point.	`set interfaces interface apply-macro juniper-port-profile Wireless Access Point`
Set VLAN on VLANs stanza.	`set vlans vlan namevlan-id vlan-id`
Set port family to **ethernet-ewitching** Set port mode to **Access**.	`set interfaces interface unit 0 family ethernet-switching port-mode access`
Set VLAN on port stanza.	`set interfaces interface unit 0 family ethernet-switching vlan members vlan-members`
Set RSTP protocol with edge option.	`set protocols rstp interface interface edge`
RSTP protocol is disabled if redundant trunk groups are configured.	`set protocols rstp interface interface disable`
Routed Uplink Port Role	

Table 125: Port Role Configuration Summary *(continued)*

Configuration Description	CLI Commands
Set the port role to Routed Uplink.	`set interfaces interface apply-macro juniper-port-profile Routed Uplink`
Set port family to inet. Set IP address on the port.	`set interfaces interfaceunit 0 family inet address ipaddress`

Set class-of-service parameters
 	
SCHEDULER_MAP=juniper-port-profile-map
 	
IEEE_CLASSIFIER=juniper-ieee-classifier
 	
DSCP_CLASSIFIER=juniper-dscp-classifier
	`set class-of-service interfaces interfacescheduler-map juniper-port-profile-map`
	`set class-of-service interfaces interface unit 0 classifiers ieee-802.1 juniper_ieee_classifier`
	`set class-of-service interfaces interfaceunit 0 classifiers dscp juniper-dscp-classifier`

Set CoS configuration	Refer to Table 126 on page 918 for details.
Layer 2 Uplink Port Role	
Set the port role to **Layer 2 Uplink**.	`set interfaces interface apply-macro juniper-port-profile Layer2 Uplink`
Set port family to **ethernet-switching** Set port mode to **trunk**.	`set interfaces interface unit 0 family ethernet-switching port-mode trunk`
Set Native VLAN name.	`set interfaces interface unit 0 family ethernet-switching native-vlan-id vlan-name`
Set the port as part of all valid VLANs; "valid" refers to all VLANs except native VLAN and voice VLANs.	`set interfaces interface unit 0 family ethernet-switching vlan members vlan-members`
Set port security parameter.	`set ethernet-switching-options secure-access-port dhcp-trusted`
Set RSTP protocol with point-to-point option.	`set protocols rstp interface interface mode point-to-point`
Disable RSTP if redundant trunk groups are configured.	`set protocols rstp interface interface disable`

Set class-of-service parameters.
 	
SCHEDULER_MAP=juniper-port-profile-map
 	
IEEE_CLASSIFIER=juniper_ieee_classifier
 	
DSCP_CLASSIFIER=juniper_dscp_classifier
	`set class-of-service interfaces interfacescheduler-map juniper-port-profile-map`
	`set class-of-service interfaces interface unit 0 classifiers ieee-802.1 juniper_ieee_classifier`
	`set class-of-service interfaces interfaceunit 0 classifiers dscp juniper-dscp-classifier`

Set CoS configuration	Refer to Table 126 on page 918 for details.

Table 126 on page 918 lists the CLI commands for the recommended CoS settings that are committed when the CoS configuration is set.

Table 126: Recommended CoS Settings for Port Roles

CoS Parameter	CLI Command
Forwarding Classes	
voice	`set class-of-service forwarding-classes class voice queue-num 7`
expedited-forwarding	`set class-of-service forwarding-classes class expedited-forwarding queue-num 5`
assured-forwarding	`set class-of-service forwarding-classes class assured-forwarding queue-num 1`
best-effort	`set class-of-service forwarding-classes class best-effort queue-num 0`
Schedulers	

strict-priority-scheduler	The CLI commands are:
	
	`set class-of-service schedulers strict-priority-scheduler transmit-rate percent 10`
	`set class-of-service schedulers strict-priority-scheduler buffer-size percent 5`
	`set class-of-service schedulers strict-priority-scheduler priority strict-high`

expedited-scheduler	The CLI commands are:
	
	`set class-of-service schedulers expedited-scheduler transmit-rate percent 30`
	`set class-of-service schedulers expedited-scheduler buffer-size percent 30`
	`set class-of-service schedulers expedited-scheduler priority low`

assured-scheduler	The CLI commands are:
	
	`set class-of-service schedulers assured-scheduler transmit-rate percent 25`
	`set class-of-service schedulers strict-priority-scheduler buffer-size percent 25`
	`set class-of-service schedulers strict-priority-scheduler priority low`

best-effort-scheduler	The CLI commands are:
	
	`set class-of-service schedulers best-effort-scheduler transmit-rate percent 35`
	`set class-of-service schedulers best-effort-scheduler buffer-size percent 40`
	`set class-of-service schedulers best-effort-scheduler priority low`

Classifiers	The classifiers are:
	
	`set class-of-service classifiers ieee-802.1 juniper_ieee_classifier import default forwarding-class voice loss-priority low code-points 101`
	`set class-of-service classifiers dscp juniper_dscp_classifier import default forwarding-class voice loss-priority low code-points 101110`

Related Documentation

- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 909
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 919

Configuring Gigabit Ethernet Interfaces (CLI Procedure)

An Ethernet interface must be configured for optimal performance in a high-traffic network. J-EX Series switches include a factory default configuration that:

- Enables all the network interfaces on the switch
- Sets a default port mode (access)
- Sets default link settings
- Specifies a logical unit (**unit 0**) and assigns it to **family ethernet-switching** (except on J-EX8200 switches)
- Specifies Rapid Spanning Tree Protocol (RSTP) and Link Layer Discovery Protocol (LLDP)

This topic describes:

- Configuring VLAN Options and Port Mode on page 919
- Configuring the Link Settings on page 919
- Configuring the IP Options on page 920

Configuring VLAN Options and Port Mode

The factory default configuration includes a default VLAN and enables interfaces for the access port mode. Access interfaces typically connect to network devices such as PCs, printers, IP telephones, and IP cameras.

If you are connecting a desktop phone or wireless access point or a security camera to a Power over Ethernet (PoE) port, you can configure some parameters for the PoE interface. The PoE interfaces are enabled by default. For detailed information on PoE settings, see “Configuring PoE (CLI Procedure)” on page 3021.

If you are connecting a device to other switches and to routers on the LAN, you need to assign the interface to a logical port and configure the logical port as a trunk port. See “Port Role Configuration with the J-Web Interface (with CLI References)” on page 915 for more information about port configuration.

To configure a Gigabit Ethernet interface or 10-Gigabit Ethernet interface for trunk port mode:

```
[edit]
user@switch# set interfaces interface-name unit logical-unit-number family ethernet-switching
port-mode trunk
```

Configuring the Link Settings

J-EX Series switches include a factory default configuration that enables interfaces with the following link settings:

- All Gigabit Ethernet interfaces are set to **auto-negotiation**.
- The speed for Gigabit Ethernet interfaces is set to **auto**, allowing the interface to operate at 10m, 100m or 1g. The link operates at the highest possible speed, depending on the capabilities of the remote end.
- The flow control for Gigabit Ethernet interfaces and 10-Gigabit Ethernet interfaces is set to **enabled**.
- The link mode is set to **auto**, allowing the interface to operate as either full duplex or half duplex. The link operates as full duplex unless this mode is not supported at the remote end.
- The 10-Gigabit Ethernet interfaces default to **no auto-negotiation**. The default speed is 10g and the default link mode is full duplex.

To configure the link settings:

- Set link settings for a Gigabit Ethernet interface:

```
[edit]
user@switch# set interfaces ge-fpc/pic/port ether-options
```

- Set link settings for a 10-Gigabit Ethernet interface:

```
[edit]
user@switch# set interfaces xe-fpc/1/port ether-options
```

NOTE: An uplink port in a J-EX4200 switch always has a PIC value of 1.

For a J-EX4200 standalone switch, *fpc* refers to the switch itself and is always 0. In a Virtual Chassis configuration, *fpc* refers to the member ID. In a J-EX8200 switch, *fpc* refers to the line card number.

The **ether-options** statement allows you to modify the configuration:

- **802.3ad**—Specify an aggregated Ethernet bundle. See "Configuring Aggregated Ethernet Interfaces (CLI Procedure)" on page 922.
- **auto-negotiation**—Enable or disable autonegotation of flow control, link mode, and speed.
- **flow-control**—Enable or disable flow control.
- **link-mode**—Specify **full-duplex**, **half-duplex**, or **automatic**.
- **speed**—Specify **10m**, **100m**, **1g**, or **autonegotiation**.

Configuring the IP Options

To specify an IP address for the logical unit using IPv4:

```
[edit]
user@switch# set interfaces interface-name unit logical-unit-number family inet address ip-address
```

To specify an IP address for the logical unit using IPv6:

```
[edit]
user@switch# set interfaces interface-name unit logical-unit-number family inet6 address
ip-address
```

NOTE: Access interfaces on J-EX4200 switches are set to family ethernet-switching by default. You might have to delete this or another user-configured family setting before changing the setting to family inet or family inet6.

Related Documentation

- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 909
- Monitoring Interface Status and Traffic on page 931
- **show interfaces ge- on page 1005**
- **show interfaces xe- on page 1019**
- Understanding Interface Naming Conventions on J-EX Series Switches on page 865

Setting the Mode on an SFP+ Uplink Module (CLI Procedure)

SFP+ uplink modules are supported on J-EX4200 switches. You can use these uplink modules either for two SFP+ transceivers or four SFP transceivers. You configure the operating mode on the module to match the type of transceiver you want to use—that is, for SFP+ transceivers, you configure the 10-gigabit operating mode, and for SFP transceivers, you configure the 1-gigabit operating mode.

By default, the SFP+ uplink module operates in the 10-gigabit mode and supports only SFP+ transceivers. If you have not changed the module from the default setting and you want to use SFP+ transceivers, you do not need to configure the operating mode.

To set the operating mode of an SFP+ uplink module, change the operating mode to the appropriate mode for the transceiver type you want to use by using one of the following commands:

```
[edit]
user@switch# set chassis fpc 0 pic 1 sfpplus pic-mode 1g
```

```
[edit]
user@switch# set chassis fpc 0 pic 1 sfpplus pic-mode 10g
```

The changed operating mode takes effect immediately unless a port on the SFP+ uplink module is a Virtual Chassis port (VCP). If any port on the SFP+ uplink module is a VCP, the changed operating mode does not take effect until the next reboot of the switch.

NOTE: During the operating mode change, the Packet Forwarding Engine is restarted. In a Virtual Chassis configuration, this means that the Flexible PIC Concentrator connection with the master is dropped and then reconnected.

You can see whether the operating mode has been changed to the new mode you configured by issuing the **show chassis pic fpc-slot *slot-number* pic-slot 1** command.

Related Documentation

- Uplink Modules in J-EX4200 Switches
- Optical Interface Support in J-EX4200 Switches

Configuring Aggregated Ethernet Interfaces (CLI Procedure)

Use the link aggregation feature to aggregate one or more links to form a virtual link or link aggregation group (LAG). The MAC client can treat this virtual link as if it were a single link. Link aggregation increases bandwidth, provides graceful degradation as failure occurs, and increases availability.

NOTE: An interface with an already configured IP address cannot form part of the aggregation group.

To configure aggregated Ethernet interfaces, using the CLI:

1. Specify the number of aggregated Ethernet interfaces to be created:

```
[edit chassis]
user@switch# set aggregated-devices ethernet device-count 2
```

2. Specify the minimum number of links for the aggregated Ethernet interface (ae*x*), that is, the defined bundle, to be labeled “up”:

NOTE: By default only one link must be up for the bundle to be labeled “up”.

```
[edit interfaces]
user@switch# set ae0 aggregated-ether-options minimum-links 2
```

3. Specify the link speed for the aggregated Ethernet bundle:

```
[edit interfaces]
user@switch# set ae0 aggregated-ether-options link-speed 10g
```

4. Specify the members to be included within the aggregated Ethernet bundle:

```
[edit interfaces]
user@switch# set xe-0/1/0 ether-options 802.3ad ae0
user@switch# set xe-1/1/0 ether-options 802.3ad ae0
```

5. Specify an interface family for the aggregated Ethernet bundle:

```
[edit interfaces]
user@switch# set ae0 unit 0 family inet address 192.0.2.0/25
```

For information about adding LACP to a LAG, see “Configuring Aggregated Ethernet LACP (CLI Procedure)” on page 926.

Related Documentation

- Configuring Aggregated Ethernet Interfaces (J-Web Procedure) on page 923
- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 740

- Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 746
- Verifying the Status of a LAG Interface on page 932
- Understanding Aggregated Ethernet Interfaces and LACP on page 867

Configuring Aggregated Ethernet Interfaces (J-Web Procedure)

Use the link aggregation feature to aggregate one or more Ethernet interfaces to form a virtual link or link aggregation group (LAG) on a J-EX Series switch. The MAC client can treat this virtual link as if it were a single link. Link aggregation increases bandwidth, provides graceful degradation as failure occurs, and increases availability. You can use the J-Web interface to configure aggregated Ethernet interfaces, or a LAG, on the switch.

NOTE: Interfaces that are already configured with MTU, duplex, flow control, or logical interfaces are listed but are not available for aggregation.

To configure an aggregated Ethernet interface (also referred to as a LAG):

1. Select **Configure** > **Interfaces** > **Link Aggregation**.

   The list of aggregated interfaces is displayed.

   

   NOTE: After you make changes to the configuration in this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options** > **Commit**. See "Using the Commit Options to Commit Configuration Changes (J-Web Procedure)" on page 334 for details about all commit options.

2. Click one of the following:
   - **Add**—Creates an aggregated Ethernet interface, or LAG. Enter information as specified in Table 127 on page 924.
   - **Edit**—Modifies a selected LAG.
     - **Aggregation**—Modifies settings for the selected LAG. Enter information as specified in Table 127 on page 924.
     - **VLAN**—Specifies VLAN options for the selected LAG. Enter information as specified in Table 128 on page 925.
     - **IP Option**—Specifies IP options for the selected LAG. Enter information as specified in Table 129 on page 925.
   - **Delete**—Deletes the selected LAG.

- **Disable Port** or **Enable Port**—Disables or enables the administrative status on the selected interface.
- **Device Count**—Configures the number of aggregated logical devices available to the switch. Select the number and click **OK**.

Table 127: Aggregated Ethernet Interface Options

Field	Function	Your Action
Aggregated Interface	Specifies the name of the aggregated interface.	None. The name is supplied by the software.
LACP Mode	Specifies the mode in which LACP packets are exchanged between the interfaces. The modes are: • **None**—Indicates that no mode is applicable. • **Active**—Indicates that the interface initiates transmission of LACP packets • **Passive**—Indicates that the interface responds only to LACP packets.	Select from the list.
Description	Specifies a description for the LAG.	Enter a description.
Interface	Specifies the interfaces in the LAG.	To add interfaces to the LAG, select the interfaces and click **Add**. Click **OK**. To remove an interface from the LAG, select the interface and click **Remove**. NOTE: Only interfaces that are configured with the same speed can be selected together for a LAG.
Enable Log	Specifies whether to enable generation of log entries for the LAG.	Select the check box to enable log generation, or clear the check box to disable log generation.

Table 128: VLAN Options

Field	Function	Your Action
Port Mode	Specifies the mode of operation for the port: trunk or access.	If you select **Trunk**, you can: 1. Click **Add** to add a VLAN member. 2. Select the VLAN and click **OK**. 3. (Optional) Associate a native VLAN ID with the port. If you select **Access**, you can: 1. Select the VLAN member to be associated with the port. 2. (Optional) Associate a VoIP VLAN with the interface. Only a VLAN with a VLAN ID can be associated as a VoIP VLAN. Click **OK**.

Table 129: IP Options

Field	Function	Your Action
IPv4 Address	Specifies an IPv4 address for the selected LAG.	1. Select the check box **IPv4 address**. 2. Type an IP address—for example, **10.10.10.10**. 3. Enter the subnet mask or address prefix. For example, 24 bits represents **255.255.255.0**. 4. Click **OK**.
IPv6 Address	Specifies an IPv6 address for the selected LAG.	1. Select the check box **IPv6 address**. 2. Type an IP address—for example, **2001:ab8:85a3::8a2e:370:7334**. 3. Enter the subnet mask or address prefix. 4. Click **OK**.

Related Documentation

- Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 922
- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 740
- Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 746
- Verifying the Status of a LAG Interface on page 932
- Configuring Aggregated Ethernet LACP (CLI Procedure) on page 926
- Understanding Aggregated Ethernet Interfaces and LACP on page 867

Configuring Aggregated Ethernet LACP (CLI Procedure)

For aggregated Ethernet interfaces on J-EX Series switches, you can configure the Link Aggregation Control Protocol (LACP). LACP is one method of bundling several physical interfaces to form one logical interface. You can configure aggregated Ethernet with or without LACP enabled.

Before you configure LACP, be sure you have:

- Configured the aggregated Ethernet bundles—also known as link aggregation groups (LAGs). See "Configuring Aggregated Ethernet Interfaces (CLI Procedure)" on page 922

When LACP is enabled, the local and remote sides of the aggregated Ethernet links exchange protocol data units (PDUs), containing information about the state of the link. You can configure Ethernet links to actively transmit PDUs, or you can configure the links to passively transmit them, sending out LACP PDUs only when they receive them from another link. One side of the link must be configured as **active** for the link to be up.

NOTE: Do not add LACP to a LAG if the remote end of the LAG link is a security device, unless the security device supports LACP. Security devices often do not support LACP because they require a deterministic configuration.

To configure LACP:

1. Enable one side of the aggregated Ethernet link as active:

```
[edit interfaces]
user@switch# set aex aggregated-ether-options lacp active
```

2. Specify the interval at which the interfaces send LACP packets:

```
[edit interfaces]
user@switch# set aex aggregated-ether-options lacp periodic fast
```

Related Documentation

- Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 922
- Configuring Aggregated Ethernet Interfaces (J-Web Procedure) on page 923
- Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 746
- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 740
- Verifying the Status of a LAG Interface on page 932
- Understanding Aggregated Ethernet Interfaces and LACP on page 867

Configuring Unicast RPF (CLI Procedure)

Unicast reverse-path forwarding (RPF) can help protect your LAN from denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks on untrusted interfaces. Enabling unicast RPF on the switch interfaces filters traffic with source addresses that do not use the incoming interface as the best return path back to the source. When a packet comes into an interface, if that interface is not the best return path to the source, the switch discards the packet. If the incoming interface is the best return path to the source, the switch forwards the packet.

NOTE: On J-EX4200 switches, you can only enable unicast RPF globally, on all switch interfaces. You cannot enable unicast RPF on a per-interface basis.

Before you begin:

- On a J-EX8200 switch, ensure that the selected switch interface is symmetrically routed before you enable unicast RPF. A symmetrically routed interface is an interface that uses the same route in both directions between the source and the destination. Do not enable unicast RPF on asymmetrically routed interfaces. An asymmetrically routed interface uses different paths to send and receive packets between the source and the destination.
- On a J-EX4200 switch, ensure that *all* switch interfaces are symmetrically routed before you enable unicast RPF on an interface. When you enable unicast RPF on any interface, it is enabled globally on all switch interfaces. Do not enable unicast RPF on asymmetrically routed interfaces. An asymmetrically routed interface uses different paths to send and receive packets between the source and the destination.

To enable unicast RPF, configure it explicitly on a selected customer-edge interface:

```
[edit interfaces]
user@switch# set ge-1/0/10 unit 0 family inet rpf-check
```

BEST PRACTICE: On J-EX4200 switches, unicast RPF is enabled globally on *all* switch interfaces, regardless of whether you configure it explicitly on only one interface or only on some interfaces.

On J-EX4200 switches, we recommend that you enable unicast RPF explicitly on either all interfaces or only one interface. To avoid possible confusion, do not enable it on only some interfaces:

- Enabling unicast RPF explicitly on only one interface makes it easier if you choose to disable it in the future because you must explicitly disable unicast RPF on every interface on which you explicitly enabled it. If you explicitly enable unicast RPF on two interfaces and you disable it on only one interface, unicast RPF is still implicitly enabled globally on the switch. The drawback to this approach is that the switch displays the flag that indicates

that unicast RPF is enabled only on interfaces on which unicast RPF is explicitly enabled, so even though unicast RPF is enabled on all interfaces, this status is not displayed.

- Enabling unicast RPF explicitly on all interfaces makes it easier to know whether unicast RPF is enabled on the switch because every interface shows the correct status. (Only interfaces on which you explicitly enable unicast RPF display the flag that indicates that unicast RPF is enabled.) The drawback to this approach is that if you want to disable unicast RPF, you must explicitly disable it on every interface. If unicast RPF is enabled on any interface, it is implicitly enabled on all interfaces.

Related Documentation

- Example: Configuring Unicast RPF on a J-EX Series Switch on page 900
- Verifying Unicast RPF Status on page 935
- Disabling Unicast RPF (CLI Procedure) on page 928
- Troubleshooting Unicast RPF on page 941
- Understanding Unicast RPF for J-EX Series Switches on page 872

Disabling Unicast RPF (CLI Procedure)

Unicast reverse-path forwarding (RPF) can help protect your LAN from denial-of-service (DoS) and distributed denial-of-service (DDoS) attacks on untrusted interfaces. Unicast RPF filters traffic with source addresses that do not use the incoming interface as the best return path back to the source. If the network configuration changes so that an interface that has unicast RPF enabled becomes a trusted interface or becomes asymmetrically routed (the interface that receives a packet is not the best return path to the packet's source), disable unicast RPF.

To disable unicast RPF on a J-EX4200 switch, you must delete it from every interface on which you explicitly configured it. If you do not disable unicast RPF on every interface on which you explicitly enabled it, it remains implicitly enabled on all interfaces. If you attempt to delete unicast RPF from an interface on which it was not explicitly enabled, the message **warning: statement not found** displays. If you do not disable unicast RPF on every interface on which you explicitly enabled it, unicast RPF remains implicitly enabled on all interfaces of the J-EX4200 switch.

On J-EX8200 switches, the switch does not apply unicast RPF to an interface unless you explicitly enable that interface for unicast RPF.

To disable unicast RPF, delete its configuration from the interface:

```
[edit interfaces]
user@switch# delete ge-1/0/10 unit 0 family inet rpf-check
```

NOTE: On J-EX4200 switches, if you do not disable unicast RPF on every interface on which you explicitly enabled it, unicast RPF remains implicitly enabled on all interfaces.

Related Documentation

- Example: Configuring Unicast RPF on a J-EX Series Switch on page 900
- Verifying Unicast RPF Status on page 935
- Configuring Unicast RPF (CLI Procedure) on page 927
- Understanding Unicast RPF for J-EX Series Switches on page 872

Configuring IP Directed Broadcast (CLI Procedure)

You can use IP directed broadcast on a J-EX Series switch to facilitate remote network management by sending broadcast packets to hosts on a specified subnet without broadcasting to the entire network. IP directed broadcast packets are broadcast on only the target subnet. The rest of the network treats IP directed broadcast packets as unicast packets and forwards them accordingly.

Before you begin to configure IP directed broadcast:

- Ensure that the subnet on which you want broadcast packets using IP direct broadcast is not directly connected to the Internet.
- Configure a routed VLAN interface (RVI) for the subnet that will be enabled for IP direct broadcast. See “Configuring Routed VLAN Interfaces (CLI Procedure)” on page 1137 or “Configuring VLANs for J-EX Series Switches (J-Web Procedure)” on page 1133.

NOTE: We recommend that you do not enable IP directed broadcast on subnets that have a direct connection to the Internet because of increased exposure to denial-of-service (DoS) attacks.

To enable IP directed broadcast for a specified subnet:

1. Add the target subnet’s logical interfaces to the VLAN:

```
[edit interfaces]
user@switch# set ge-0/0/0.0 family ethernet-switching vlan members v1

user@switch# set ge-0/0/1.0 family ethernet-switching vlan members v1
```

2. Configure the Layer 3 interface on the VLAN that is the target of the IP directed broadcast packets:

```
[edit interfaces]
user@switch# set vlan.1 family inet address 10.1.2.1/24
```

3. Associate a Layer 3 interface with the VLAN:

```
[edit vlans]
user@switch# set v1 l3-interface vlan.1
```

4. Enable the Layer 3 interface for the VLAN to receive IP directed broadcasts:

```
[edit interfaces]
user@switch# set vlan.1 family inet targeted-broadcast
```

Related Documentation

- Example: Configuring IP Directed Broadcast on a J-EX Series Switch on page 904
- Understanding IP Directed Broadcast for J-EX Series Switches on page 876

Configuring a Layer 3 Subinterface (CLI Procedure)

J-EX Series switches use Layer 3 subinterfaces to divide a physical interface into multiple logical interfaces, each corresponding to a VLAN. The switch uses the Layer 3 subinterfaces to route traffic between subnets.

To configure Layer 3 subinterfaces, you enable VLAN tagging and partition one or more physical ports into multiple logical interfaces, each corresponding to a VLAN ID.

Before you begin, make sure you set up your VLANs. See “Configuring VLANs for J-EX Series Switches (CLI Procedure)” on page 1136 or “Configuring VLANs for J-EX Series Switches (J-Web Procedure)” on page 1133.

To configure Layer 3 subinterfaces:

1. Enable VLAN tagging:

```
[edit interfaces interface-name]
user@switch# set vlan-tagging
```

2. Bind each VLAN ID to a logical interface:

```
[edit interfaces interface-name]
user@switch# set unit logical-unit-number vlan-id vlan-id-number
```

Related Documentation

- Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch on page 893
- Verifying That Layer 3 Subinterfaces Are Working on page 934
- Understanding Layer 3 Subinterfaces on page 871

CHAPTER 53

Verifying Interfaces

Monitoring Interface Status and Traffic

Purpose Use the monitoring functionality to view interface status or to monitor interface bandwidth utilization and traffic statistics on the J-EX Series switches.

The J-Web interface monitors interface bandwidth utilization and plots real-time charts to display input and output rates in bytes per second. In addition, the Interface monitoring page displays input and output packet counters and error counters in the form of charts.

Alternatively, you can enter the show commands in the CLI to view interface status and traffic statistics.

Action To view general interface information in the J-Web interface such as available interfaces, select **Monitor > Interfaces**. Click any interface to view details about its status.

To set up interface monitoring for Virtual Chassis and J-EX8200 switches, select a member from the **Port for FPC** list. Details such as the admin status and link status are displayed in the table.

NOTE: By default, the details of the first member in the Port for FPC drop-down list is displayed.

You have the following options:

- **Start/Stop**—Starts or stops monitoring the selected interface.
- **Show Graph**—Displays input and output packet counters and error counters in the form of charts. Also, click on the pop-up icon to view the graph in a separate window.

- **Details**—Displays interface information such as general details, traffic statistics, I/O errors, CoS counters, and Ethernet statistics.
- **Refresh Interval (sec)**—Displays the time interval you have set for page refresh.
- **Clear Statistics**—Clears the statistics for the interface selected from the table.

Using the CLI:

- To view interface status for all the interfaces, enter **show interfaces xe-**.
- To view status and statistics for a specific interface, enter **show interfaces xe-*interface-name***.
- To view status and traffic statistics for all interfaces, enter either **show interfaces xe- detail** or **show interfaces xe- extensive**.

Meaning In the J-Web interface the charts displayed are:

- Bar charts—Display the input and output error counters.
- Pie charts—Display the number of broadcast, unicast, and multicast packet counters.

For details about output from the CLI commands, see **show interfaces ge-** (Gigabit Ethernet) or **show interfaces xe-** (10-Gigabit Ethernet).

Related Documentation
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 909
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 919

Verifying the Status of a LAG Interface

Purpose Verify that a LAG (**ae0**) has been created on the switch.

Action **show interfaces aeo terse**

```
Interface             Admin   Link  Proto      Local             Remote

ae0                    up       up

ae0.0                  up       up     inet     10.10.10.2/24
```

Meaning The output confirms that the **ae0** link is up and shows the **family** and IP address assigned to this link.

Related Documentation
- Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 922
- Configuring Aggregated Ethernet Interfaces (J-Web Procedure) on page 923

- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 740

Verifying That LACP Is Configured Correctly and Bundle Members Are Exchanging LACP Protocol Packets

To verify that LACP has been set up correctly and that the bundle members are transmitting LACP protocol packets.

1. Verifying the LACP Setup on page 933
2. Verifying That the LACP Packets Are Being Exchanged on page 933

Verifying the LACP Setup

Purpose Verify that the LACP has been set up correctly.

Action Use the **show lacp interfaces *interface-name*** command to check that LACP has been enabled as active on one end.

show lacp interfaces xe-0/1/0

```
show lacp interfaces xe-0/1/0

Aggregated interface: ae0

    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity

      xe-0/1/0      Actor    No   Yes    No   No   No   Yes     Fast    Active

      xe-0/1/0    Partner    No   Yes    No   No   No   Yes     Fast   Passive

    LACP protocol:   Receive State    Transmit State          Mux State

      xe-0/1/0          Defaulted     Fast periodic           Detached
```

Meaning This example shows that LACP has been configured with one side as **active** and the other as **passive**. When LACP is enabled, one side must be set as **active** in order for the bundled link to be **up**.

Verifying That the LACP Packets Are Being Exchanged

Purpose Verify that LACP packets are being exchanged between interfaces.

Action Use the **show interfaces ae*x* statistics** command to display LACP BPDU exchange information.

show interfaces ae0 statistics

```
Physical interface: ae0, Enabled, Physical link is Down
  Interface index: 153, SNMP ifIndex: 30
  Link-level type: Ethernet, MTU: 1514, Speed: Unspecified, Loopback: Disabled,
```

```
Source filtering: Disabled, Flow control: Disabled, Minimum links needed: 1,
Minimum bandwidth needed: 0
Device flags   : Present Running
Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
Current address: 02:19:e2:50:45:e0, Hardware address: 02:19:e2:50:45:e0
Last flapped   : Never
Statistics last cleared: Never
  Input packets : 0
  Output packets: 0
Input errors: 0, Output errors: 0

Logical interface ae0.0 (Index 71) (SNMP ifIndex 34)
  Flags: Hardware-Down Device-Down SNMP-Traps Encapsulation: ENET2
  Statistics        Packets        pps         Bytes          bps
  Bundle:
      Input :             0          0             0            0
      Output:             0          0             0            0
  Protocol inet,
    Flags: None
    Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
      Destination: 10.10.10/24, Local: 10.10.10.1, Broadcast: 10.10.10.255
```

Meaning The output here shows that the link is down and that no PDUs are being exchanged (when there is no other traffic flowing on the link).

Related Documentation

- Configuring Aggregated Ethernet LACP (CLI Procedure) on page 926
- Verifying the Status of a LAG Interface on page 932
- Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 746

Verifying That Layer 3 Subinterfaces Are Working

Purpose After configuring Layer 3 subinterfaces, verify they are set up properly and transmitting data.

Action

1. Use the **show interfaces** command to determine if you successfully created the subinterfaces and the links are up:

   ```
   user@switch> show interfaces interface-name terse

   Interface               Admin Link Proto    Local                 Remote
   ge-0/0/0                up    up
   ge-0/0/0.0              up    up   inet     1.1.1.1/24
   ge-0/0/0.1              up    up   inet     2.1.1.1/24
   ge-0/0/0.2              up    up   inet     3.1.1.1/24
   ge-0/0/0.3              up    up   inet     4.1.1.1/24
   ge-0/0/0.4              up    up   inet     5.1.1.1/24
   ge-0/0/0.32767          up    up
   ```

2. Use the **ping** command from a device on one subnet to an address on another subnet to determine if packets were transmitted correctly on the subinterface VLANs:

   ```
   user@switch> ping ip-address
   ```

```
PING 1.1.1.1 (1.1.1.1): 56 data bytes
64 bytes from 1.1.1.1: icmp_seq=0 ttl=64 time=0.157 ms
64 bytes from 1.1.1.1: icmp_seq=1 ttl=64 time=0.238 ms
64 bytes from 1.1.1.1: icmp_seq=2 ttl=64 time=0.255 ms
64 bytes from 1.1.1.1: icmp_seq=3 ttl=64 time=0.128 ms
--- 1.1.1.1 ping statistics ---
4 packets transmitted, 4 packets received, 0% packet loss
```

Meaning The output confirms that the subinterfaces are created and the links are up.

Related Documentation

- Configuring a Layer 3 Subinterface (CLI Procedure) on page 930
- Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch on page 893

Verifying Unicast RPF Status

Purpose Verify that unicast reverse-path forwarding (RPF) is enabled and is working on the interface.

Action Use one of the **show interfaces *interface-name*** commands with either the **extensive** or **detail** options to verify that unicast RPF is enabled and working on the switch. The example below displays output from the **show interfaces ge- extensive** command.

```
user@switch> show interfaces ge-1/0/10 extensive
Physical interface: ge-1/0/10, Enabled, Physical link is Down
  Interface index: 139, SNMP ifIndex: 58, Generation: 140
  Link-level type: Ethernet, MTU: 1514, Speed: Auto, MAC-REWRITE Error: None,
  Loopback: Disabled, Source filtering: Disabled, Flow control: Enabled,
  Auto-negotiation: Enabled, Remote fault: Online
  Device flags   : Present Running
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:19:e2:50:95:ab, Hardware address: 00:19:e2:50:95:ab
  Last flapped   : Never
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :                   0                   0 bps
   Output bytes  :                   0                   0 bps
   Input  packets:                   0                   0 pps
   Output packets:                   0                   0 pps
   IPv6 transit statistics:
    Input  bytes  :                  0
    Output bytes  :                  0
    Input  packets:                  0
    Output packets:                  0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
    FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

    FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
```

```
Queue counters:       Queued packets  Transmitted packets      Dropped packets

  0 best-effort                    0                    0                    0

  1 assured-forw                   0                    0                    0

  5 expedited-fo                   0                    0                    0

  7 network-cont                   0                    0                    0

Active alarms  : LINK
Active defects : LINK
MAC statistics:                       Receive         Transmit
  Total octets                              0                0
  Total packets                             0                0
  Unicast packets                           0                0
  Broadcast packets                         0                0
  Multicast packets                         0                0
  CRC/Align errors                          0                0
  FIFO errors                               0                0
  MAC control frames                        0                0
  MAC pause frames                          0                0
  Oversized frames                          0
  Jabber frames                             0
  Fragment frames                           0
  VLAN tagged frames                        0
  Code violations                           0
Filter statistics:
  Input packet count                        0
  Input packet rejects                      0
  Input DA rejects                          0
  Input SA rejects                          0
  Output packet count                                        0
  Output packet pad count                                    0
  Output packet error count                                  0
  CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Incomplete
Packet Forwarding Engine configuration:
  Destination slot: 1

Logical interface ge-1/0/10.0 (Index 69) (SNMP ifIndex 59) (Generation 135)
  Flags: Device-Down SNMP-Traps 0x0 Encapsulation: ENET2
  Traffic statistics:
   Input  bytes  :                   0
   Output bytes  :                   0
   Input  packets:                   0
   Output packets:                   0
   IPv6 transit statistics:
    Input  bytes  :                  0
    Output bytes  :                  0
    Input  packets:                  0
    Output packets:                  0
  Local statistics:
   Input  bytes  :                   0
   Output bytes  :                   0
   Input  packets:                   0
   Output packets:                   0
  Transit statistics:
   Input  bytes  :                   0                    0 bps
   Output bytes  :                   0                    0 bps
```

```
Input  packets:                 0                    0 pps
Output packets:                 0                    0 pps
IPv6 transit statistics:
 Input  bytes  :                0
 Output bytes  :                0
 Input  packets:                0
 Output packets:                0
   Protocol inet, Generation: 144, Route table: 0
 Flags: uRPF
 Addresses, Flags: Is-Preferred Is-Primary
```

Meaning The **show interfaces ge-1/0/10 extensive** command (and the **show interfaces ge-1/0/10 detail** command) displays in-depth information about the interface. The **Flags:** output field near the bottom of the display reports the unicast RPF status. If unicast RPF has not been enabled, the **uRPF** flag is not displayed.

On J-EX4200 switches, unicast RPF is implicitly enabled on *all* switch interfaces, including aggregated Ethernet interfaces (also referred to as link aggregation groups or LAGs) and routed VLAN interfaces (RVIs) when you enable unicast RPF on a single interface. However, the unicast RPF status is shown as enabled only on interfaces for which you have explicitly configured unicast RPF. Thus, the **uRPF** flag is not displayed on interfaces for which you have not explicitly configured unicast RPF even though unicast RPF is implicitly enabled on all interfaces on J-EX4200 switches.

Related Documentation

- **show interfaces xe- on page 1019**
- Example: Configuring Unicast RPF on a J-EX Series Switch on page 900
- Configuring Unicast RPF (CLI Procedure) on page 927
- Disabling Unicast RPF (CLI Procedure) on page 928
- Troubleshooting Unicast RPF on page 941

Verifying IP Directed Broadcast Status

Purpose Verify that IP directed broadcast is enabled and is working on the subnet.

Action Use the **show vlans extensive** command to verify that IP directed broadcast is enabled and working on the subnet as shown in the following example.

Related Documentation

- Configuring IP Directed Broadcast (CLI Procedure) on page 929
- Example: Configuring IP Directed Broadcast on a J-EX Series Switch on page 904

CHAPTER 54

Troubleshooting Interfaces

Troubleshooting Network Interfaces on J-EX4200 Switches

This topic provides troubleshooting information for specific problems related to interfaces on J-EX4200 switches.

The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP+ uplink module is down

Problem The interface on the port in which an SFP or SFP+ transceiver is installed in an SFP+ uplink module installed in a J-EX4200 switch is down.

When you check the status with the CLI command **show interfaces ge-** or with the J-Web user interface, the disabled port is not listed.

Cause By default, the SFP+ uplink module operates in the 10-gigabit mode and supports only SFP+ transceivers. The operating mode for the module is incorrectly set.

Solution Either SFP+ or SFP transceivers can be installed in SFP+ uplink modules. You must configure the operating mode of the SFP+ uplink module to match the type of transceiver you want to use. For SFP+ transceivers, configure the 10-gigabit operating mode and for SFP transceivers, configure the 1-gigabit operating mode. See “Setting the Mode on an SFP+ Uplink Module (CLI Procedure)” on page 921.

Related Documentation

- Troubleshooting Uplink Module Installation or Replacement on J-EX4200 Switches on page 942
- Monitoring Interface Status and Traffic on page 931

- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 919
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 909
- Removing a Transceiver from a J-EX Series Switch
- Uplink Modules in J-EX4200 Switches
- J-EX Series Switches Interfaces Overview on page 863

Troubleshooting an Aggregated Ethernet Interface

Problem The **show interfaces terse** command shows that the LAG is down.

Solution Check the following:

- Verify that there is no configuration mismatch.
- Verify that all member ports are up.
- Verify that a LAG is part of family ethernet—switching (Layer 2 LAG) or family inet (Layer 3 LAG).
- Verify that the LAG member is connected to the correct LAG at the other end.
- Verify that the LAG members belong to the same switch (or the same Virtual Chassis).

Related Documentation

- Verifying the Status of a LAG Interface on page 932
- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 740
- Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 746

Troubleshooting Interface Configuration and Cable Faults

Troubleshooting interface configuration and connectivity on the J-EX Series switch:

1. Interface Configuration or Connectivity Is Not Working on page 940

Interface Configuration or Connectivity Is Not Working

Problem You encounter errors when you attempt to configure an interface on the switch, or the interface is exhibiting connectivity problems.

Solution Use the port troubleshooter feature in the J-Web interface to identify and rectify port configuration and connectivity related problems.

To use the J-Web interface port troubleshooter:

1. Select the option **Troubleshoot** from the main menu.
2. Click **Troubleshoot Port**. The Port Troubleshooting wizard is displayed. Click **Next**.

3. Select the ports to troubleshoot.
4. Select the test cases to be executed on the selected port. Click **Next**.

   When the selected test cases are executed, the final result and the recommended action is displayed.

If there is a cable fault, the port troubleshooter displays details and the recommended action. For example, the cable must be replaced.

If the port configuration needs to be modified, the port troubleshooter displays details and the recommended action.

Related Documentation

- Monitoring Interface Status and Traffic on page 931
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 909
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 919
- Connecting and Configuring a J-EX Series Switch (CLI Procedure) on page 161
- Connecting and Configuring a J-EX Series Switch (J-Web Procedure) on page 163

Troubleshooting Unicast RPF

Troubleshooting issues for unicast reverse-path forwarding (RPF) on J-EX Series switches include:

1. Legitimate Packets Are Discarded on page 941

Legitimate Packets Are Discarded

Problem The switch filters valid packets from legitimate sources, which results in the switch's discarding packets that should be forwarded.

Solution The interface or interfaces on which legitimate packets are discarded are asymmetrically routed interfaces. An asymmetrically routed interface uses different paths to send and receive packets between the source and the destination, so the interface that receives a packet is not the same interface the switch uses to reply to the packet's source.

Unicast RPF works properly only on symmetrically routed interfaces. A symmetrically routed interface is an interface that uses the same route in both directions between the source and the destination. Unicast RPF filters packets by checking the forwarding table for the best return path to the source of an incoming packet. If the best return path uses the same interface as the interface that received the packet, the switch forwards the packet. If the best return path uses a different interface than the interface that received the packet, the switch discards the packet.

NOTE: On J-EX4200 switches, unicast RPF works properly only if all switch interfaces—including aggregated Ethernet interfaces (also referred to as link aggregation groups or LAGs) and routed VLAN interfaces (RVIs)—are symmetrically routed, because unicast RPF is enabled globally on all switch interfaces.

Related Documentation

- Verifying Unicast RPF Status on page 935
- Understanding Unicast RPF for J-EX Series Switches on page 872

Troubleshooting Uplink Module Installation or Replacement on J-EX4200 Switches

This topic provides troubleshooting information for specific problems related to uplink module ports on J-EX4200 switches.

1. Virtual Chassis port (VCP) connection does not work on page 942

Virtual Chassis port (VCP) connection does not work

Problem The Virtual Chassis port (VCP) connection configured in a J-EX4200 switch does not work.

A port of the uplink module is set as a VCP.

Cause The uplink module installed in the switch was replaced.

Solution Set a port in the uplink module as a VCP. See “Setting an Uplink Module Port as a Virtual Chassis Port (CLI Procedure)” on page 792.

Related Documentation

- Monitoring Interface Status and Traffic on page 931
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 919
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 909
- Installing an Uplink Module in a J-EX4200 Switch
- Removing a Transceiver from a J-EX Series Switch
- Uplink Modules in J-EX4200 Switches
- Understanding Virtual Chassis Hardware Configuration on a J-EX4200 Switch

CHAPTER 55

Configuration Statements for Interfaces

[edit chassis] Configuration Statement Hierarchy

```
chassis {
  aggregated-devices {
    ethernet {
      device-count  number;
    }
  }
  auto-image-upgrade;
  }
  fpc slot {
    pic pic-number {
      sfpplus {
        pic-mode mode;
    power-budget-priority priority;
    }
  lcd-menu fpc slot-number {
    menu-item (menu-name | menu-option);
  }
}
  psu {
    redundancy {
      n-plus-n;
    }
  }
  redundancy{
    graceful-switchover ;
  }
```

Related Documentation

- Understanding Aggregated Ethernet Interfaces and LACP on page 867
- Understanding Power Management on J-EX Series Switches on page 302

[edit interfaces] Configuration Statement Hierarchy

```
interfaces {
  aex {
```

```
    aggregated-ether-options {
      (flow-control | no-flow-control);
      lacp mode {
        periodic interval;
      }
      link-speed speed;
      minimum-links number;
    }
    description text;
    disable;
    hold-time up milliseconds down milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;
    traceoptions;
    (traps | no-traps);
    unit logical-unit-number {
      description text;
      disable;
      family family-name {...}
      proxy-arp (restricted | unrestricted);
      (traps | no-traps);
      vlan-id vlan-id-number;
    }
    vlan-tagging;
  }
  fe-fpc/pic/port {
    description text;
    disable;
    mtu bytes;
    no-gratuitous-arp-request;
    speed speed;
    traceoptions;
    (traps | no-traps);
    unit logical-unit-number {
      description text;
      disable;
      family family-name {...}
      proxy-arp (restricted | unrestricted);
      (traps | no-traps);
      vlan-id vlan-id-number;
    }
    vlan-tagging;
  }
  ge-fpc/pic/port {
    description text;
    disable;
    ether-options {
      802.3ad aex {
        lacp {
          force-up;
        }
      }
      (auto-negotiation | no-auto-negotiation);
      (flow-control | no-flow-control);
      link-mode mode;
      speed (auto-negotiation | speed);
```

```
    }
    hold-time up milliseconds down milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;
    traceoptions;
    (traps | no-traps);
    unit logical-unit-number {
      description text;
      disable;
      family family-name {...}
      proxy-arp (restricted | unrestricted);
      rpm;
      (traps | no-traps);
      vlan-id vlan-id-number;
    }
    vlan-tagging;
  }
  interface-range  interface-range name {
    description text;
    disable;
    ether-options {
      802.3ad aex {
        lacp {
          force-up;
        }
      }
      (auto-negotiation | no-auto-negotiation);
      (flow-control| no-flow-control);
      link-mode mode;
      speed (auto-negotiation | speed);
    }
    hold-time up milliseconds down milliseconds;
    member interface-name;
    member-range  starting-interface name to  ending-interface name;
    mtu bytes;
    unit logical-unit-number {
      description text;
      disable;
      family family-name {...}
      proxy-arp (restricted | unrestricted);
      rpm;
      (traps | no-traps);
      vlan-id vlan-id-number;
    }
  }
  lo0 {
    description text;
    disable;
    hold-time up milliseconds down milliseconds;
    traceoptions;
    (traps | no-traps);
    unit logical-unit-number {
      description text;
      disable;
      family family-name {...}
      (traps | no-traps);
```

```
    }
  }
  me0 {
    description text;
    disable;
    hold-time up milliseconds down milliseconds;
    no-gratuitous-arp-request;
    traceoptions;
    (traps | no-traps);
    unit logical-unit-number {
      description text;
      disable;
      family family-name {...}
      (traps | no-traps);
      vlan-id vlan-id-number;
    }
    vlan-tagging;
  }
  vlan {
    description  text;
    disable;
    hold-time up milliseconds down milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;
    traceoptions;
    (traps | no-traps);
    unit logical-unit-number {
      description text;
      disable;
      family family-name {...}
      proxy-arp (restricted | unrestricted);
      (traps | no-traps);
    }
  }
  vme {
    description text;
    disable;
    hold-time up milliseconds down milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;
    traceoptions;
    (traps | no-traps);
    unit logical-unit-number {
      description text;
      disable;
      family family-name {...}
      (traps | no-traps);
      vlan-id vlan-id-number;
    }
    vlan-tagging;
  }
  xe-fpc/pic/port {
    description text;
    disable;
    ether-options {
      802.3ad aex {
```

```
        lacp (802.3ad) {
          force-up;
        }
      }
      (auto-negotiation | no-auto-negotiation);
      (flow-control | no-flow-control);
      link-mode mode;
      speed (auto-negotiation | speed);
    }
    hold-time up milliseconds down milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;
    traceoptions;
    (traps | no-traps);
    unit logical-unit-number {
      description text;
      disable;
      family family-name {...}
      proxy-arp (restricted | unrestricted);
      rpm;
      (traps | no-traps);
      vlan-id vlan-id-number;
    }
    vlan-tagging;
  }
}
```

Related Documentation

- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 919
- Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 922
- Configuring a Layer 3 Subinterface (CLI Procedure) on page 930
- Configuring Routed VLAN Interfaces (CLI Procedure) on page 1137
- Configuring the Virtual Management Ethernet Interface for Global Management of a Virtual Chassis (CLI Procedure) on page 797
- J-EX Series Switches Interfaces Overview on page 863
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

802.3ad

Syntax

```
802.3ad aex {
  lacp {
    force-up;
  }
```

Hierarchy Level [edit **interfaces** *interface-name* **ether-options**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the aggregated Ethernet logical interface number.

Options **ae***x*—Aggregated Ethernet logical interface number.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 740
- Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 746
- Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 922
- Configuring Aggregated Ethernet LACP (CLI Procedure) on page 926
- Understanding Aggregated Ethernet Interfaces and LACP on page 867
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

aggregated-devices

Syntax

```
aggregated-devices {
  ethernet {
    device-count number;
  }
}
```

Hierarchy Level [edit chassis]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure properties for aggregated devices on the switch.

The statements are explained separately.

Default Disabled.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 740
- Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 922
- Understanding Aggregated Ethernet Interfaces and LACP on page 867
- *Junos OS System Basics Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

aggregated-ether-options

Syntax

```
aggregated-ether-options {
  (flow-control | no-flow-control);
  lacp {
    (active | passive);
    admin-key key;
    periodic interval;
    system-id mac-address;
  }
  (link-protection | no-link-protection);
  link-speed speed;
  (loopback | no-loopback);
  minimum-links number;
}
```

Hierarchy Level [edit **interfaces (for EX Series switches)** ae*x*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the aggregated Ethernet properties of a specific aggregated Ethernet interface.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 740
- Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 746
- Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 922
- Configuring Aggregated Ethernet LACP (CLI Procedure) on page 926
- Understanding Aggregated Ethernet Interfaces and LACP on page 867
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

auto-negotiation

Syntax (auto-negotiation | no-auto-negotiation);

Hierarchy Level [edit **interfaces** *interface-name* **ether-options**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Explicitly enable or disable autonegotiation.

- **auto-negotiation**—Enable autonegotiation.
- **no-auto-negotiation**—Disable autonegotiation. When autonegotiation is disabled, you must explicitly configure link mode and speed options.

Default Autonegotiation is automatically enabled. No explicit action is taken after the autonegotiation is complete or if the negotiation fails.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 919
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 909
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

chassis

Syntax

```
chassis {
  aggregated-devices {
    ethernet {
      device-count number;
    }
  }
  auto-image-upgrade;
  fpc slot {
    pic pic-number {
      sfpplus {
        pic-modemode;
      }
    }
    power-budget-priority priority;
  }
  lcd-menu fpc slot-number {
    menu-item (menu-name | menu-option);
  }
  nssu {
    upgrade-group group-name {
      fpcs (slot-number | [list-of-slot-numbers]);
      member member-id {
        fpcs (slot-number | [list-of-slot-numbers]);
      }
    }
  }
  psu {
    redundancy (Power Management) {
      n-plus-n;
    }
  }
  redundancy {
    graceful-switchover;
  }
}
```

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure chassis-specific properties.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 922
- Upgrading Software Using Automatic Software Download on J-EX Series Switches on page 82

- Configuring Graceful Routing Engine Switchover in a Virtual Chassis Configuration (CLI Procedure) on page 801
- Configuring the Power Priority of Line Cards (CLI Procedure) on page 308
- Configuring Power Supply Redundancy (CLI Procedure) on page 307
- *Junos OS System Basics Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

description

Syntax description *text*;

Hierarchy Level [edit **interfaces** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Provide a textual description of the interface or the logical unit. Any descriptive text you include is displayed in the output of the show interfaces commands and is also exposed in the **ifAlias** Management Information Base (MIB) object. It has no effect on the operation of the interface or the switch.

Default No textual description is configured.

Options ***text***—Text to describe the interface. If the text includes spaces, enclose the entire text in straight quotation marks.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 919
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 909
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

device-count

Syntax device-count *number*;

Hierarchy Level [edit **chassis aggregated-devices ethernet**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the number of aggregated Ethernet logical devices available to the switch.

Default There is no default. You must configure a value.

Options ***number***—Maximum number of aggregated Ethernet logical interfaces on the switch.

Range: 1 through 64 for J-EX4200 switches
1 through 255 for J-EX8200 switches

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 740
- Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 922
- *Junos OS System Basics Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

ether-options

Syntax

```
ether-options {
  802.3ad aex {
    lacp {
      force-up;
    }
  }
  auto-negotiation;
  flow-control;
  link-mode mode;
  speed (speed | auto-negotiation);
}
```

Hierarchy Level [edit **interfaces** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure **ether-options** properties for a Gigabit Ethernet interface on the J-EX Series switch.

The remaining statements are explained separately.

Default Enabled.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 919
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 909
- Understanding Aggregated Ethernet Interfaces and LACP on page 867
- J-EX Series Switches Interfaces Overview on page 863
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

ethernet

Syntax

```
ethernet {
  device-count number;
}
```

Hierarchy Level [edit **chassis aggregated-devices**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure properties for Ethernet aggregated devices on the switch.

The remaining statement is explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 922
- *Junos OS System Basics Configuration Guide* at http://www.juniper.net/techpubs/software/junos/index.html

family (for J-EX Series switches)

Syntax

family ccc

```
family ccc;
```

family ethernet-switching

```
family ethernet-switching {
  filter input filter-name;
  filter output filter-name;
  native-vlan-id vlan-id;
  port-mode mode;
  vlan {
    members [ (all | names | vlan-ids) ];
  }
}
```

family inet

```
family inet {
  address address {
    primary;
    vrrp-group group-id {
      advertise-interval milliseconds;
      preempt | no-preempt {
        hold-time seconds;
      }
      priority number;
      virtual-address [addresses];
      virtual-link-local-address ip-address;
    }
  }
  filter input filter-name;
  filter output filter-name;
  primary;
  rpf-check;
  targeted-broadcast;
}
```

family inet6

```
family inet6 {
  address address {
    primary;
    vrrp-inet6-group group-id {
      inet6-advertise-interval milliseconds;
      preempt | no-preempt {
        hold-time seconds;
      }
      priority number;
      virtual-inet6-address [addresses];
      virtual-link-local-address ipv6–address;
    }
  }
```

```
    filter input filter-name;
    filter output filter-name;
    policer input policer-name;
    policer output policer-name;
    rpf-check;
}
```

family iso

```
family iso {
    address interface-address;
}
```

family mpls

```
family mpls;
```

Hierarchy Level

[edit **interfaces** *interface-name* **unit** *logical-unit-number*]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure protocol family information for the logical interface on the switch.

Most standard Junos OS configuration statements are available in Junos OS for J-EX Series switches. This topic lists standard Junos OS statements that you commonly use when configuring protocol families for interfaces on J-EX Series switches as well as statements that are used to configure protocol families only on switch interfaces. For information about additional standard Junos OS statements that you can configure on interfaces, see the *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.

Default

Access interfaces on J-EX4200 switches are set to **family ethernet-switching** by default. If you are going to change the family setting for an interface, you might have to delete this default setting or any user-configured family setting before you change the setting to another family type.

J-EX8200 switch interfaces do not have a default family setting.

You must configure a logical interface to be able to use the physical device.

Options See Table 130 on page 959 for protocol families available on the switch interfaces. Different protocol families support different subsets of the interfaces types on the switch.

Interface types on the switch are:

- Aggregated Ethernet (**ae**)
- Gigabit Ethernet (**ge**)
- Interface-range configuration (**interface-range**)
- Loopback (**lo0**)
- Management Ethernet (**me0**)
- Routed VLAN interface (RVI) (**vlan**)
- Virtual management Ethernet (**vme**)
- 10-Gigabit Ethernet (**xe**)

If you are using an interface range, the supported protocol families are the ones supported by the interface types that compose the range.

Not all interface types support all **family** substatements. Check your switch CLI for supported substatements for a particular protocol family configuration.

Table 130: Protocol Families and Supported Interface Types

Family	Description	Supported Interface Types						
		ae	ge	lo0	me0	vlan	vme	xe
ccc	Circuit cross-connect protocol family	✓	✓					✓
ethernet-switching	Ethernet switching protocol family	✓	✓		✓			✓
inet	IPv4 protocol family	✓	✓	✓	✓	✓	✓	✓
inet6	IPv6 protocol family	✓	✓	✓	✓	✓	✓	✓
iso	Junos OS protocol family for IS-IS traffic	✓	✓	✓	✓	✓	✓	✓
mpls	MPLS protocol family	✓	✓	✓	✓		✓	✓

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring MPLS on J-EX Series Switches on page 3071
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 919
- Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 922
- Configuring Routed VLAN Interfaces (CLI Procedure) on page 1137
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

filter

Syntax filter (input | output) *filter-name*;

Hierarchy Level [edit **interfaces** *interface-name* **unit** *logical-unit-number* **family** ethernet-switching],
[edit **interfaces** *interface-name* **unit** *logical-unit-number* **family** inet],
[edit **interfaces** *interface-name* **unit** *logical-unit-number* **family** inet6]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply a firewall filter to traffic entering the port or Layer 3 interface or exiting the Layer 3 interface.

Default All incoming traffic is accepted unmodified on the port or Layer 3 interface, and all outgoing traffic is sent unmodified from the port or Layer 3 interface.

Options ***filter-name***—Name of a firewall filter defined in the **filter** statement.

- **input**—Apply a firewall filter to traffic entering the port or Layer 3 interface.
- **output**—Apply a firewall filter to traffic exiting the Layer 3 interface.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 919
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 909
- Configuring Firewall Filters (CLI Procedure) on page 2779
- Configuring Firewall Filters (J-Web Procedure) on page 2784
- Firewall Filters for J-EX Series Switches Overview on page 2721
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

flow-control

Syntax (flow-control | no-flow-control);

Hierarchy Level [edit **interfaces** *interface-name* **ether-options**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Explicitly enable flow control, which regulates the flow of packets from the switch to the remote side of the connection, or disable it.

- **flow-control**—Enable flow control; flow control is useful when the remote device is a Gigabit Ethernet switch.
- **no-flow-control**—Disable flow control.

Default Flow control enabled.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 919
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 909
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

force-up

Syntax force-up;

Hierarchy Level [edit **interfaces** *interface-name* **ether-options 802.3ad lacp**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set the state of the interface as UP when the peer has limited LACP capability.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 919
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 909
- Understanding Aggregated Ethernet Interfaces and LACP on page 867
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

interface-range

Syntax

```
interface-range interface-range name {
  ether-options {
    802.3ad aex ;
    auto-negotiation;
    flow-control;
    link-mode mode;
    speed ( speed | auto-negotiation) ;
  }
  hold-time up milliseconds down milliseconds;
  member interface-name;
  member-range starting-interface name to ending-interface name;
  mtu bytes;
}
```

Hierarchy Level [edit interfaces]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Group interfaces that share a common configuration profile.

NOTE: The interface range definition is supported only for Gigabit, 10-Gigabit, and Fast Ethernet interfaces.

Options *interface-range-name*—Name of the interface range.

NOTE: You can use regular expressions and wildcards to specify the interfaces in the member-range configuration. Do not use wildcards for interface types.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 919
- Understanding Interface Ranges on J-EX Series Switches on page 869
- J-EX Series Switches Interfaces Overview on page 863
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

interfaces (for J-EX Series switches)

Syntax

interfaces ae

```
aex {
  aggregated-ether-options {
    (flow-control | no-flow-control);
    lacp mode {
      periodic interval;
    }
    link-speed speed;
    minimum-links number;
  }
  description text;
  disable;
  hold-time up milliseconds down milliseconds;
  mtu bytes;
  no-gratuitous-arp-request;
  traceoptions;
  (traps | no-traps);
  unit logical-unit-number {
    description text;
    disable;
    family family-name {...}
    proxy-arp (restricted | unrestricted);
    (traps | no-traps);
    vlan-id vlan-id-number;
  }
  vlan-tagging;
}
```

interfaces ge

```
ge-fpc/pic/port {
  description text;
  disable;
  ether-options {
    802.3ad aex {
      lacp {
        force-up;
      }
    }
    (auto-negotiation | no-auto-negotiation);
    (flow-control | no-flow-control);
    link-mode mode;
    speed (auto-negotiation | speed);
  }
  hold-time up milliseconds down milliseconds;
  mtu bytes;
  no-gratuitous-arp-request;
```

```
    traceoptions;
    (traps | no-traps);
    unit logical-unit-number {
      description text;
      disable;
      family family-name {...}
      proxy-arp (restricted | unrestricted);
      rpm;
      (traps | no-traps);
      vlan-id vlan-id-number;
    }
    vlan-tagging;
  }
```

interfaces interface-range

```
  interface-range interface-range name {
    description text;
    disable;
    ether-options {
      802.3ad aex {
        lacp {
          force-up;
        }
      }
      (auto-negotiation | no-auto-negotiation);
      (flow-control| no-flow-control);
      link-mode mode;
      speed (auto-negotiation | speed);
    }
    hold-time up milliseconds down milliseconds;
    member interface-name;
    member-range starting-interface name to ending-interface name;
    mtu bytes;
    unit logical-unit-number {
      description text;
      disable;
      family family-name {...}
      proxy-arp (restricted | unrestricted);
      rpm;
      (traps | no-traps);
      vlan-id vlan-id-number;
    }
  }
```

interfaces lo0

```
lo0 {
  description text;
  disable;
  hold-time up milliseconds down milliseconds;
  traceoptions;
  (traps | no-traps);
  unit logical-unit-number {
    description text;
    disable;
    family family-name {...}
    (traps | no-traps);
  }
}
```

interfaces me0

```
me0 {
  description text;
  disable;
  hold-time up milliseconds down milliseconds;
  no-gratuitous-arp-request;
  traceoptions;
  (traps | no-traps);
  unit logical-unit-number {
    description text;
    disable;
    family family-name {...}
    (traps | no-traps);
    vlan-id vlan-id-number;
  }
  vlan-tagging;
}
```

interfaces vlan

```
vlan {
  description  text;
  disable;
  hold-time up milliseconds down milliseconds;
  mtu bytes;
  no-gratuitous-arp-request;
  traceoptions;
  (traps | no-traps);
  unit logical-unit-number {
    description text;
    disable;
    family family-name {...}
    proxy-arp (restricted | unrestricted);
    (traps | no-traps);
  }
}
```

interfaces vme

```
vme {
  description text;
  disable;
  hold-time up milliseconds down milliseconds;
  mtu bytes;
  no-gratuitous-arp-request;
  traceoptions;
  (traps | no-traps);
  unit logical-unit-number {
    description text;
    disable;
    family family-name {...}
    (traps | no-traps);
    vlan-id vlan-id-number;
  }
  vlan-tagging;
}
```

interfaces xe

```
xe-fpc/pic/port {
  description text;
  disable;
  ether-options {
    802.3ad aex {
      lacp (802.3ad) {
        force-up;
      }
    }
    (auto-negotiation | no-auto-negotiation);
    (flow-control | no-flow-control);
    link-mode mode;
    speed (auto-negotiation | speed);
  }
  hold-time up milliseconds down milliseconds;
  mtu bytes;
  no-gratuitous-arp-request;
  traceoptions;
  (traps | no-traps);
  unit logical-unit-number {
    description text;
    disable;
    family family-name {...}
    proxy-arp (restricted | unrestricted);
    rpm;
    (traps | no-traps);
    vlan-id vlan-id-number;
  }
  vlan-tagging;
}
```

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure interfaces on J-EX Series switches.

Most standard Junos OS configuration statements are available in Junos OS for J-EX Series switches. This topic lists standard Junos OS statements that you commonly use when configuring interfaces on J-EX Series switches as well as statements that are used to configure only switch interfaces; it does not list all of the possible interface configuration statements for each interface.

For information about additional standard Junos OS statements that you can configure on interfaces, see the *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.

Options See Table 131 on page 967 for the interface types and protocol-family options supported on the switch. Different protocol families support different subsets of the interface types on the switch. See the **family** statement for syntax of the protocol families supported for switch interfaces.

Not all interface types support all **family** substatements. Check your switch CLI for supported substatements for a particular protocol family configuration.

Table 131: Interface Types and Their Supported Protocol Families

Interface Typ	Description	Supported Protocol Families					
		ccc	ethernet-switching	inet	inet6	iso	mpls
ae	Aggregated Ethernet interface (also referred to as a link aggregation group [LAG])	✓	✓	✓	✓	✓	✓
ge	Gigabit Ethernet interface	✓	✓	✓	✓	✓	✓
lo0	Loopback interface			✓	✓	✓	✓
me0	Management Ethernet interface		✓	✓	✓	✓	✓
vlan	Routed VLAN interface (RVI)			✓	✓	✓	
vme	Virtual management Ethernet interface			✓	✓	✓	✓
xe	10-Gigabit Ethernet interface	✓	✓	✓	✓	✓	✓
interface-range	Interface-range configuration	Supported protocol families are the ones supported by the interface types that compose the range.					

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 919
- Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 922
- Configuring a Layer 3 Subinterface (CLI Procedure) on page 930
- Configuring Routed VLAN Interfaces (CLI Procedure) on page 1137
- Configuring the Virtual Management Ethernet Interface for Global Management of a Virtual Chassis (CLI Procedure) on page 797
- J-EX Series Switches Interfaces Overview on page 863
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

lacp

Syntax

```
lacp mode {
  periodic interval;
}
```

Hierarchy Level [edit **interfaces** ae*x* **aggregated-ether-options**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the Link Aggregation Control Protocol (LACP).

Default LACP is not enabled.

Options *mode* —LACP mode:

- **active**—Initiate transmission of LACP packets
- **passive**—Respond to LACP packets

The remaining statement is explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 746
- Configuring Aggregated Ethernet LACP (CLI Procedure) on page 926
- Configuring Aggregated Ethernet Interfaces (J-Web Procedure) on page 923
- Understanding Aggregated Ethernet Interfaces and LACP on page 867
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

lacp (802.3ad)

Syntax

```
lacp {
  force-up;
}
```

Hierarchy Level [edit **interfaces** *interface-name* **ether-options 802.3ad**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the Link Aggregation Control Protocol (LACP) parameters for interfaces.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 740
- Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 746
- Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 922
- Configuring Aggregated Ethernet LACP (CLI Procedure) on page 926
- Understanding Aggregated Ethernet Interfaces and LACP on page 867
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

link-mode

Syntax link-mode *mode*;

Hierarchy Level [edit **interfaces** *interface-name* **ether-options**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set the device's link-connection characteristic.

Default The **automatic** mode is enabled.

Options *mode* —Link characteristic:

- **full-duplex**—Connection is full duplex.
- **half-duplex**—Connection is half duplex.
- **automatic**—Link mode is negotiated.

If **no-auto-negotiation** is specified in ether-options, you can select only **full-duplex** or **half-duplex**. If **auto-negotiation** is specified in **ether-options**, you can select any mode.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 919
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 909
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

link-speed

Syntax link-speed *speed*;

Hierarchy Level [edit interfaces ae*x* **aggregated-ether-options**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For aggregated Ethernet interfaces only, set the required link speed.

Options ***speed***—For aggregated Ethernet links, specify ***speed*** in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Aggregated Ethernet links on J-EX Series switches can have one of the following speed values:

- **1g**—Links are 1 Gbps.
- **10g**—Links are 10 Gbps.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 740
- Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 922

member

Syntax member *interface-name*;

Hierarchy Level [edit **interfaces interface-range** *interface-range-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the name of the member interface belonging to an interface range on the J-EX Series switch.

Options ***interface-name***—Name of the interface.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 919
- Understanding Interface Ranges on J-EX Series Switches on page 869
- J-EX Series Switches Interfaces Overview on page 863
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

members

Syntax members [(all | *names* | *vlan-ids*)];

Hierarchy Level [edit **interfaces** *interface-name* **unit** *logical-unit-number* **family ethernet-switching vlan**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For trunk interfaces, configure the VLANs for which the interface can carry traffic.

TIP: To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type ? after vlan or vlans in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.

Options **all**—Specifies that this trunk interface is a member of all the VLANs that are configured on this switch. When a new VLAN is configured on the switch, this trunk interface automatically becomes a member of the VLAN.

NOTE: Each VLAN that is configured must have a specified VLAN ID when you attempt to commit the configuration; otherwise, the configuration commit fails. Also, all cannot be the name of a VLAN on the switch.

names—Name of one or more VLANs.

vlan-ids—Numeric identifier of one or more VLANs. For a series of tagged VLANs, specify a range; for example, **10-20** or **10-20 23 27-30**.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- **show ethernet-switching interfaces on page 997**
- **show vlans on page 1263**
- Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch on page 1063
- Example: Connecting an Access Switch to a Distribution Switch on page 1078
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 919
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 909
- Creating a Series of Tagged VLANs (CLI Procedure) on page 1140
- Understanding Bridging and VLANs on J-EX Series Switches on page 1041
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

member-range

Syntax member-range *starting-interface-name* to *ending-interface-name*;

Hierarchy Level [edit **interfaces interface-range** *interface-range-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the names of the first and last members of a sequence of interfaces belonging to an interface range.

Options **Range:** *Starting interface-name* to *ending interface-name*—The name of the first member and the name of the last member in the interface sequence.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 919
- Understanding Interface Ranges on J-EX Series Switches on page 869
- J-EX Series Switches Interfaces Overview on page 863
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

minimum-links

Syntax minimum-links *number*;

Hierarchy Level [edit interfaces ae*x* **aggregated-ether-options**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For an aggregated Ethernet interface, set the minimum number of links that must be up for the bundle to be labeled up.

Options *number*—Number of links.

Range: 1 through 8 for J-EX Series switches other than J-EX8200 switches
1 through 12 for J-EX8200 switches

Default: 1

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 740
- Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 922

mtu

Syntax mtu *bytes*;

Hierarchy Level [edit **interfaces** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the maximum transmission unit (MTU) size for the media. Changing the media MTU size causes an interface to be deleted and added again. Keep the following points in mind if you are configuring MTU size for jumbo frames on these special types of interfaces:

- **For LAG interfaces**—Configuring the jumbo MTU size on a link aggregation group (LAG) interface (**ae***x*) automatically configures the jumbo MTU size on the member links.
- **For RVIs**—Jumbo frames of up to 9216 bytes are supported on the routed VLAN interface (RVI), which is named **vlan**. The RVI functions as a logical router. To route jumbo data packets on the RVI, you must configure the jumbo MTU size on the member physical interfaces of the RVI and not on the RVI itself (the **vlan** interface). However, for jumbo control packets—for example, to ping the RVI with a packet size of 6000 bytes or more—you must explicitly configure the jumbo MTU size on the interface named **vlan** (the RVI).

CAUTION: Setting or deleting the jumbo MTU size on the RVI (the vlan interface) while the switch is transmitting packets might result in dropped packets.

Default 1514 bytes

Options *bytes*—MTU size.
Range: 256 through 9216 bytes
Default: 1514 bytes

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 919
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 909
- Configuring Routed VLAN Interfaces (CLI Procedure) on page 1137
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

native-vlan-id

Syntax native-vlan-id *vlan-id*;

Hierarchy Level [edit **interfaces** *interface-name* unit 0 **family** ethernet-switching]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the VLAN identifier to associate with untagged packets received on the interface.

Options ***vlan-id***—Numeric identifier of the VLAN.
Range: 0 through 4095

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show vlans on page 1263**
- **show ethernet-switching interfaces on page 997**
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 919
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 909
- Understanding Bridging and VLANs on J-EX Series Switches on page 1041
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

periodic

Syntax periodic *interval*;

Hierarchy Level [edit **interfaces** ae*x* **aggregated-ether-options lacp**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the interval for periodic transmission of LACP packets.

Default **fast**

Options ***interval*** —Interval at which to periodically transmit LACP packets:

- **fast**—Transmit packets every second. This is the default.
- **slow**—Transmit packets every 30 seconds.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 746
- Configuring Aggregated Ethernet LACP (CLI Procedure) on page 926
- Understanding Aggregated Ethernet Interfaces and LACP on page 867
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

pic

Syntax

```
pic pic-number {
  sfpplus {
    pic-modemode;
      }
    }
```

Hierarchy Level [edit **chassis fpc** *slot*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable the specified port of the SFP+ uplink module to perform in the operating mode specified by **pic-mode**. The port is indicated by a Physical Interface Card (PIC) number.

Options ***pic-number***—Number of the PIC. For uplink ports in J-EX4200 switches, the PIC number is always 1.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Setting the Mode on an SFP+ Uplink Module (CLI Procedure) on page 921

pic-mode

Syntax pic-mode *mode*;

Hierarchy Level [edit **chassis fpc** *slot* **pic** *pic-number* **sfpplus**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the operating mode for the specified port on the SFP+ uplink module on a J-EX4200 switch.

Options ***mode***—Operating mode of the SFP+ uplink module:

- **1G**—1-gigabit operating mode
- **10G**—10-gigabit operating mode

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Setting the Mode on an SFP+ Uplink Module (CLI Procedure) on page 921

port-mode

Syntax port-mode *mode*;

Hierarchy Level [edit **interfaces** *interface-name* **unit** *logical-unit-number* **family** ethernet-switching]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure whether an interface on the switch operates in access or trunk mode.

Default All switch interfaces are in access mode.

Options

access—Have the interface operate in access mode. In this mode, the interface can be in a single VLAN only. Access interfaces typically connect to network devices such as PCs, printers, IP telephones, and IP cameras.

trunk—Have the interface operate in trunk mode. In this mode, the interface can be in multiple VLANs and can multiplex traffic between different VLANs. Trunk interfaces typically connect to other switches and to routers on the LAN.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Connecting an Access Switch to a Distribution Switch on page 1078
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 919
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 909
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

rpf-check

Syntax rpf-check;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet],
[edit interfaces *interface-name* unit *logical-unit-number* family inet6]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description On J-EX4200 switches, enable a reverse-path forwarding (RPF) check on unicast traffic (except ECMP packets) on *all* ingress interfaces.

On J-EX8200 switches, enable an RPF check on unicast traffic, including ECMP packets, on the selected ingress interface.

Default Unicast RPF is disabled on all interfaces.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Unicast RPF on a J-EX Series Switch on page 900
- Configuring Unicast RPF (CLI Procedure) on page 927
- Disabling Unicast RPF (CLI Procedure) on page 928
- Understanding Unicast RPF for J-EX Series Switches on page 872

sfpplus

Syntax

```
sfpplus {
  pic-mode mode;
}
```

Hierarchy Level [edit **chassis fpc** *slot* **pic** *pic-number*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the operating mode for the specified port on the SFP+ uplink module on the J-EX4200 switch.

The remaining statement is explained separately.

Default By default, the SFP+ uplink module operates in the 10-gigabit mode and supports SFP+ transceivers.

NOTE: The SFP+ uplink module provides two ports for 10-gigabit small form-factor pluggable (SFP+) transceivers when configured to operate in 10-gigabit mode or four ports for 1-gigabit small form-factor pluggable (SFP) transceivers when configured to operate in 1-gigabit mode.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Setting the Mode on an SFP+ Uplink Module (CLI Procedure) on page 921

speed

Syntax speed (auto-negotiation | *speed*) ;

Hierarchy Level [edit **interfaces** *interface-name* **ether-options**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the interface's speed.

Default If the **auto-negotiation** statement at the **[edit interfaces *interface-name* ether-options]** hierarchy level is enabled, the **auto-negotiation** option is enabled by default.

Options

- **auto-negotiation**—Automatically negotiate the speed based on the speed of the other end of the link. This option is available only when the **auto-negotiation** statement at the **[edit interfaces *interface-name* ether-options]** hierarchy level is enabled.
- ***speed***—Specify the interface speed. If the **auto-negotiation** statement at the **[edit interfaces *interface-name* ether-options]** hierarchy level is disabled, you must specify a specific value. This value sets the speed that is used on the link. If the **auto-negotiation** statement is enabled, you might want to configure a specific speed value to advertise the desired speed to the remote end.
  - **10m**—10 Mbps
  - **100m**—100 Mbps
  - **1g**—1 Gbps

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 919
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 909
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

targeted-broadcast

Syntax targeted-broadcast;

Hierarchy Level [edit **interfaces**ge-*chassis/slot/port* **unit** *logical-unit-number* **family inet**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable IP directed broadcast on a specified subnet.

Default IP directed broadcast is disabled.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring IP Directed Broadcast on a J-EX Series Switch on page 904
- Configuring IP Directed Broadcast (CLI Procedure) on page 929
- Understanding IP Directed Broadcast for J-EX Series Switches on page 876

unit

Syntax

```
unit logical-unit-number {
  description text;
  disable;
  family family-name {...}
  proxy-arp (restricted | unrestricted);
  rpm;
  (traps | no-traps);
  vlan-id vlan-id-number;
}
```

Hierarchy Level [edit **interfaces** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options ***logical-unit-number***—Number of the logical unit.

Range: 0 through 16,384

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 919
- Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 922
- J-EX Series Switches Interfaces Overview on page 863
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

vlan

Syntax

```
vlan {
  members [ (all | names | vlan-ids) ];
}
```

Hierarchy Level [edit **interfaces** *interface-name* **unit** *logical-unit-number* **family** ethernet-switching]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Bind an 802.1Q VLAN tag ID to a logical interface.

The remaining statement is explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- **show ethernet-switching interfaces on page 997**
- Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches on page 1070
- Configuring Routed VLAN Interfaces (CLI Procedure) on page 1137
- Understanding Bridging and VLANs on J-EX Series Switches on page 1041
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

vlan-id

Syntax vlan-id *vlan-id-number*;

Hierarchy Level [edit **interfaces** *interface-name* **unit** *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Bind an 802.1Q VLAN tag ID to a logical interface.

NOTE: The VLAN tag ID cannot be configured on logical interface unit 0. The logical unit number must be 1 or higher.

Options ***vlan-id-number***—A valid VLAN identifier.
Range: 1 through 4094

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- **vlan-tagging on page 987**
- Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch on page 893
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 919
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 909
- Configuring a Layer 3 Subinterface (CLI Procedure) on page 930
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

vlan-tagging

Syntax vlan-tagging;

Hierarchy Level [edit **interfaces** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable VLAN tagging. The switch will receive and forward single-tag frames with 802.1Q VLAN tags.

Default VLAN tagging is disabled by default.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- **vlan-id on page 986**
- Example: Configuring Layer 3 Subinterfaces for a Distribution Switch and an Access Switch on page 893
- Configuring a Layer 3 Subinterface (CLI Procedure) on page 930
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

CHAPTER 56

Operational Mode Commands for Interfaces

clear ipv6 neighbors

Syntax
clear ipv6 neighbors
<all | host *hostname*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Clear IPv6 neighbor cache information.

Options
none—Clear all IPv6 neighbor cache information.

all—(Optional) Clear all IPv6 neighbor cache information.

host *hostname*—(Optional) Clear the information for the specified IPv6 neighbors.

Required Privilege Level
view

Related Documentation
- **show ipv6 neighbors on page 1031**

List of Sample Output
clear ipv6 neighbors on page 990

Output Fields
When you enter this command, you are provided feedback on the status of your request.

clear ipv6 neighbors
```
user@host> clear ipv6 neighbors
```

monitor interface

Syntax monitor interface
<*interface-name* | traffic <detail>>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display real-time statistics about interfaces, updating the statistics every second. Check for and display common interface failures, such as loopbacks detected and increases in framing errors.

Options none—Display real-time statistics for all interfaces.

interface-name—(Optional) Display real-time statistics for the specified interface.

traffic—(Optional) Display traffic data for all active interfaces.

detail—(Optional) With traffic option only, display detailed output.

Additional Information The output of this command shows how much each field has changed since you started the command or since you cleared the counters by using the c key. For a description of the statistical information provided in the output of this command, see the **show interfaces extensive** command for a particular interface type in the *Junos OS Interfaces Command Reference*. To control the output of the **monitor interface *interface-name*** command while it is running, use the keys listed in Table 132 on page 991. The keys are not case-sensitive.

Table 132: Output Control Keys for the monitor interface interface-name Command

Key	Action
c	Clears (returns to zero) the delta counters since **monitor interface** was started. This does not clear the accumulative counter. To clear the accumulative counter, use the **clear interfaces interval** command.
f	Freezes the display, halting the display of updated statistics and delta counters.
i	Displays information about a different interface. The command prompts you for the name of a specific interface.
n	Displays information about the next interface. The **monitor interface** command displays the physical or logical interfaces in the same order as the **show interfaces terse** command.
q or Esc	Quits the command and returns to the command prompt.
t	Thaws the display, resuming the update of the statistics and delta counters.

To control the output of the **monitor interface traffic** command while it is running, use the keys listed in Table 133 on page 992. The keys are not case-sensitive.

Table 133: Output Control Keys for the monitor interface traffic Command

Key	Action
b	Displays the statistics in units of bytes and bytes per second (bps).
c	Clears (return to 0) the delta counters in the **Current Delta** column. The statistics counters are not cleared.
d	Displays the **Current Delta** column (instead of the rate column) in bps or packets per second (pps).
p	Displays the statistics in units of packets and packets per second (pps).
q or Esc	Quits the command and returns to the command prompt.
r	Displays the rate column (instead of the **Current Delta** column) in bps and pps.

Required Privilege Level trace

List of Sample Output

monitor interface (Physical) on page 993
monitor interface (OTN Interface) on page 994
monitor interface (Logical) on page 995
monitor interface traffic on page 995
monitor interface traffic detail on page 996

Output Fields Table 134 on page 992 describes the output fields for the **monitor interface** command. Output fields are listed in the approximate order in which they appear.

Table 134: monitor interface Output Fields

Field Name	Field Description	Level of Output
router1	Hostname of the router.	All levels
Seconds	How long the monitor interface command has been running or how long since you last cleared the counters.	All levels
Time	Current time (UTC).	All levels
Delay *x/y/z*	Time difference between when the statistics were displayed and the actual clock time. • *x*—Time taken for the last polling (in milliseconds). • *y*—Minimum time taken across all pollings (in milliseconds). • *z*—Maximum time taken across all pollings (in milliseconds).	All levels
Interface	Short description of the interface, including its name, status, and encapsulation.	All levels
Link	State of the link: **Up**, **Down**, or **Test**.	All levels

Table 134: monitor interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
Current delta	Cumulative number for the counter in question since the time shown in the Seconds field, which is the time since you started the command or last cleared the counters.	All levels
Statistics	For an explanation of the interface statistics, see the description of the **show interfaces extensive** command for a particular interface type in the *Junos OS Interfaces Command Reference*.	All levels
Description	With the **traffic** option, displays the interface description configured at the **[edit interfaces *interface-name*]** hierarchy level.	**detail**

monitor interface (Physical)

```
user@host> monitor interface so-0/0/0
router1                         Seconds: 19                 Time: 15:46:29

Interface: so-0/0/0, Enabled, Link is Up
Encapsulation: PPP, Keepalives, Speed: OC48
Traffic statistics:                                        Current Delta
  Input packets:                    6045 (0 pps)                   [11]
  Input bytes:                   6290065 (0 bps)                [13882]
  Output packets:                  10376 (0 pps)                   [10]
  Output bytes:                 10365540 (0 bps)                 [9418]
Encapsulation statistics:
  Input keepalives:                 1901                           [2]
  Output keepalives:                1901                           [2]
  NCP state: Opened
  LCP state: Opened
Error statistics:
  Input errors:                        0                           [0]
  Input drops:                         0                           [0]
  Input framing errors:                0                           [0]
  Policed discards:                    0                           [0]
  L3 incompletes:                      0                           [0]
  L2 channel errors:                   0                           [0]
  L2 mismatch timeouts:                0                           [0]
  Carrier transitions:                 1                           [0]
  Output errors:                       0                           [0]
  Output drops:                        0                           [0]
  Aged packets:                        0                           [0]
Active alarms : None
Active defects: None
SONET error counts/seconds:
  LOS count                            1                           [0]
  LOF count                            1                           [0]
  SEF count                            1                           [0]
  ES-S                                 0                           [0]
  SES-S                                0                           [0]
SONET statistics:
  BIP-B1                          458871                           [0]
  BIP-B2                          460072                           [0]
  REI-L                           465610                           [0]
  BIP-B3                          458978                           [0]
  REI-P                           458773                           [0]
```

```
Received SONET overhead:
  F1      : 0x00  J0        : 0x00  K1        : 0x00
  K2      : 0x00  S1        : 0x00  C2        : 0x00
  C2(cmp) : 0x00  F2        : 0x00  Z3        : 0x00
  Z4      : 0x00  S1(cmp)   : 0x00
Transmitted SONET overhead:
  F1      : 0x00  J0        : 0x01  K1        : 0x00
  K2      : 0x00  S1        : 0x00  C2        : 0xcf
  F2      : 0x00  Z3        : 0x00  Z4        : 0x00

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'
```

monitor interface (OTN Interface)

```
user@host> monitor interface ge-7/0/0

Interface: ge-7/0/0, Enabled, Link is Up
Encapsulation: Ethernet, Speed: 10000mbps
Traffic statistics:
  Input bytes:                          0 (0 bps)
  Output bytes:                         0 (0 bps)
  Input packets:                        0 (0 pps)
  Output packets:                       0 (0 pps)
Error statistics:
  Input errors:                         0
  Input drops:                          0
  Input framing errors:                 0
  Policed discards:                     0
  L3 incompletes:                       0
  L2 channel errors:                    0
  L2 mismatch timeouts:                 0
  Carrier transitions:                  5
  Output errors:                        0
  Output drops:                         0
  Aged packets:                         0
Active alarms : None
Active defects: None
Input MAC/Filter statistics:
  Unicast packets                       0
  Broadcast packets                     0
  Multicast packets                     0
  Oversized frames                      0
  Packet reject count                   0
  DA rejects                            0
  SA rejects                            0
Output MAC/Filter Statistics:
  Unicast packets                       0
  Broadcast packets                     0
  Multicast packets                     0
  Packet pad count                      0
  Packet error count                    0
OTN Link 0
  OTN Alarms: OTU_BDI, OTU_TTIM, ODU_BDI
  OTN Defects: OTU_BDI, OTU_TTIM, ODU_BDI, ODU_TTIM
  OTN OC - Seconds
    LOS                                 2
    LOF                                 9
  OTN OTU - FEC Statistics
    Corr err ratio                    N/A
    Corr bytes                          0
    Uncorr words                        0
  OTN OTU - Counters
```

```
    BIP                                 0
    BBE                                 0
    ES                                  0
    SES                                 0
    UAS                               422
  OTN ODU - Counters
    BIP                                 0
    BBE                                 0
    ES                                  0
    SES                                 0
    UAS                               422
  OTN ODU - Received Overhead    APSPCC 0-3:              0
```

monitor interface (Logical)

```
user@host> monitor interface so-1/0/0.0
host name                Seconds: 16                  Time: 15:33:39
                                                            Delay: 0/0/1
Interface: so-1/0/0.0, Enabled, Link is Down
Flags: Hardware-Down Point-To-Point SNMP-Traps
Encapsulation: PPP
Local statistics:                                          Current delta
  Input bytes:                        0                              [0]
  Output bytes:                       0                              [0]
  Input packets:                      0                              [0]
  Output packets:                     0                              [0]
Remote statistics:
  Input bytes:                        0 (0 bps)                      [0]
  Output bytes:                       0 (0 bps)                      [0]
  Input packets:                      0 (0 pps)                      [0]
  Output packets:                     0 (0 pps)                      [0]
Traffic statistics:
  Destination address: 192.168.8.193, Local: 192.168.8.21

Next='n', Quit='q' or ESC, Freeze='f', Thaw='t', Clear='c', Interface='i'
```

monitor interface traffic

```
user@host> monitor interface traffic
host name                Seconds: 15                  Time: 12:31:09

Interface    Link  Input packets        (pps)     Output packets        (pps)
 so-1/0/0    Down              0          (0)                  0          (0)
 so-1/1/0    Down              0          (0)                  0          (0)
 so-1/1/1    Down              0          (0)                  0          (0)
 so-1/1/2    Down              0          (0)                  0          (0)
 so-1/1/3    Down              0          (0)                  0          (0)
 t3-1/2/0    Down              0          (0)                  0          (0)
 t3-1/2/1    Down              0          (0)                  0          (0)
 t3-1/2/2    Down              0          (0)                  0          (0)
 t3-1/2/3    Down              0          (0)                  0          (0)
 so-2/0/0      Up         211035          (1)              36778          (0)
 so-2/0/1      Up         192753          (1)              36782          (0)
 so-2/0/2      Up         211020          (1)              36779          (0)
 so-2/0/3      Up         211029          (1)              36776          (0)
 so-2/1/0      Up         189378          (1)              36349          (0)
 so-2/1/1    Down              0          (0)              18747          (0)
 so-2/1/2    Down              0          (0)              16078          (0)
 so-2/1/3      Up              0          (0)              80338          (0)
 at-2/3/0      Up              0          (0)                  0          (0)
 at-2/3/1    Down              0          (0)                  0          (0)

Bytes=b, Clear=c, Delta=d, Packets=p, Quit=q or ESC, Rate=r, Up=^U, Down=^D
```

monitor interface traffic detail

```
user@host> monitor interface traffic detail
host name                 Seconds: 15                 Time: 12:31:09

Interface    Link  Input packets   (pps)  Output packets   (pps)     Description

 t1-0/1/1:0  Up    19769           (0)    0                (0)       To-OSAKA-1

...

Bytes=b, Clear=c, Delta=d, Packets=p, Quit=q or ESC, Rate=r, Up=^U, Down=^D
```

show ethernet-switching interfaces

Syntax show ethernet-switching interfaces
<brief | detail | summary>
<interface *interface-name*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display information about switched Ethernet interfaces.

Options none—(Optional) Display brief information for Ethernet switching interfaces.

brief | detail | summary—(Optional) Display the specified level of output.

interface *interface-name*—(Optional) Display Ethernet switching information for a specific interface.

Required Privilege Level view

Related Documentation

- **show ethernet-switching mac-learning-log on page 1241**
- **show ethernet-switching table on page 1249**
- Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure) on page 2516

List of Sample Output

show ethernet-switching interfaces on page 998
show ethernet-switching interfaces ge-0/0/15 brief on page 999
show ethernet-switching interfaces ge-0/0/2 detail (Blocked by RTG rtggroup) on page 999
show ethernet-switching interfaces ge-0/0/15 detail (Blocked by STP) on page 999
show ethernet-switching interfaces ge-0/0/17 detail (Disabled by bpdu-control) on page 999
show ethernet-switching interfaces detail (C-VLAN to S-VLAN Mapping) on page 999

Output Fields Table 135 on page 997 lists the output fields for the **show ethernet-switching interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 135: show ethernet-switching interfaces Output Fields

Field Name	Field Description	Level of Output
Interface	Name of a switching interface.	All levels
State	Interface state. Values are **up** and **down**.	none, **brief**, **detail**, **summary**
VLAN members	Name of a VLAN.	none, **brief**, **detail**, **summary**
Tag	Number of the 802.1Q-tag.	All levels

Table 135: show ethernet-switching interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
Tagging	Specifies whether the interface forwards 802.1Q-tagged or untagged traffic.	All levels
Blocking	The forwarding state of the interface: • **unblocked**—Traffic is forwarded on the interface. • **blocked**—Traffic is not being forwarded on the interface. • **Disabled by bpdu control**—The interface is disabled due to receiving BPDUs on a protected interface. If the **disable-timeout** statement has been included in the BPDU configuration, the interface automatically returns to service after the timer expires. • **blocked by RTG**—The specified redundant trunk group is disabled. • **blocked by STP**—The interface is disabled due to a spanning tree protocol error. • **MAC limit exceeded**—The interface is temporarily disabled due to a MAC limiting error. The disabled interface is automatically restored to service when the disable timeout expires. • **MAC move limit exceeded**—The interface is temporarily disabled due to a MAC move limiting error. The disabled interface is automatically restored to service when the disable timeout expires. • **Storm control in effect**—The interface is temporarily disabled due to a storm control error. The disabled interface is automatically restored to service when the disable timeout expires.	none, **brief**, **detail**, **summary**
Index	The VLAN index internal to the Junos OS.	**detail**
mapping	The C-VLAN to S-VLAN mapping information: • **dot1q-tunneled**—The interface maps all traffic to the S-VLAN (all-in-one bundling). • **native**—The interface maps untagged and priority tagged packets to the S-VLAN. • **push**—The interface maps packets to a firewall filter to an S-VLAN. • **policy-mapped**—The interface maps packets to a specifically defined S-VLAN. • *integer*—The interface maps packets to the specified S-VLAN.	**detail**

show ethernet-switching interfaces

```
user@switch> show ethernet-switching interfaces

Interface    State  VLAN members        Tag   Tagging  Blocking

ae0.0        up     default                   untagged unblocked
ge-0/0/2.0   up     vlan300           300     untagged blocked by RTG (rtggroup)
ge-0/0/3.0   up     default                            blocked by STP
ge-0/0/4.0   down   default                            MAC limit exceeded
ge-0/0/5.0   down   default                            MAC move limit exceeded
ge-0/0/6.0   down   default                            Storm control in effect
ge-0/0/7.0   down   default                            unblocked
ge-0/0/13.0  up     default                   untagged unblocked
ge-0/0/14.0  up     vlan100             100   tagged   unblocked
                    vlan200             200   tagged   unblocked
ge-0/0/15.0  up     vlan100             100   tagged   blocked by STP
                    vlan200             200   tagged   blocked by STP
```

```
ge-0/0/16.0  down    default                  untagged unblocked
ge-0/0/17.0  down    vlan100            100   tagged   Disabled by bpdu-control

                     vlan200            200   tagged   Disabled by bpdu-control
```

show ethernet-switching interfaces ge-0/0/15 brief

```
user@switch> show ethernet-switching interfaces ge-0/0/15 brief
Interface   State  VLAN members   Tag   Tagging        Blocking

ge-0/0/15.0  up     vlan100        100   tagged         blocked by STP
                    vlan200        200   tagged         blocked by STP
```

show ethernet-switching interfaces ge-0/0/2 detail (Blocked by RTG rtggroup)

```
user@switch> show ethernet-switching interfaces ge-0/0/2 detail

Interface: ge-0/0/2.0, Index: 65, State: up, Port mode: Access
VLAN membership:
    vlan300, 802.1Q Tag: 300, untagged, msti-id: 0, blocked by RTG(rtggroup)
Number of MACs learned on IFL: 0
```

show ethernet-switching interfaces ge-0/0/15 detail (Blocked by STP)

```
user@switch> show ethernet-switching interfaces ge-0/0/15 detail

Interface: ge-0/0/15.0, Index: 70, State: up, Port mode: Trunk
VLAN membership:
    vlan100, 802.1Q Tag: 100, tagged, msti-id: 0, blocked by STP
    vlan200, 802.1Q Tag: 200, tagged, msti-id: 0, blocked by STP

Number of MACs learned on IFL: 0
```

show ethernet-switching interfaces ge-0/0/17 detail (Disabled by bpdu-control)

```
user@switch> show ethernet-switching interfaces ge-0/0/17 detail

Interface: ge-0/0/17.0, Index: 71, State: down, Port mode: Trunk
VLAN membership:
    vlan100, 802.1Q Tag: 100, tagged, msti-id: 1, Disabled by bpdu-control
    vlan200, 802.1Q Tag: 200, tagged, msti-id: 2, Disabled by bpdu-control
Number of MACs learned on IFL: 0
```

show ethernet-switching interfaces detail (C-VLAN to S-VLAN Mapping)

```
user@switch>show ethernet-switching interfaces ge-0/0/6.0 detail
Interface: ge-0/0/6.0, Index: 73, State: up, Port mode: Access
VLAN membership:
    map, 802.1Q Tag: 134, Mapped Tag: native, push, dot1q-tunneled, unblocked
    map, 802.1Q Tag: 134, Mapped Tag: 20, push, dot1q-tunneled, unblocked
```

show interfaces diagnostics optics

Syntax show interfaces diagnostics optics *interface-name*

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display diagnostics data and alarms for Gigabit Ethernet optical transceivers (SFP or SFP+) installed in J-EX Series switches. The information provided by this command is known as digital optical monitoring (DOM) information.

Thresholds that trigger a high alarm, low alarm, high warning, or low warning are set by the transponder vendors. Generally, a high alarm or low alarm indicates that the optics module is not operating properly. This information can be used to diagnose why a transceiver is not working.

Options *interface-name*—Name of the interface associated with the port in which the transceiver is installed: **ge-*fpc/pic/port*** or **xe-*fpc/pic/port***.

Required Privilege Level view

Related Documentation

- Monitoring Interface Status and Traffic on page 931
- Installing a Transceiver in a J-EX Series Switch
- Removing a Transceiver from a J-EX Series Switch
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

List of Sample Output

show interfaces diagnostics optics ge-0/1/0 (SFP Transceiver) on page 1002
show interfaces diagnostics optics xe-0/1/0 (SFP+ Transceiver) on page 1003

Output Fields Table 136 on page 1000 lists the output fields for the **show interfaces diagnostics optics** command. Output fields are listed in the approximate order in which they appear.

Table 136: show interfaces diagnostics optics Output Fields

Field Name	Field Description
Physical interface	Displays the name of the physical interface.
Laser bias current	Displays the magnitude of the laser bias power setting current, in milliamperes. The laser bias provides direct modulation of laser diodes and modulates currents.
Laser output power	Displays the laser output power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
Module temperature	Displays the temperature, in Celsius and Fahrenheit.
Module voltage	Displays the voltage, in Volts.

Table 136: show interfaces diagnostics optics Output Fields *(continued)*

Field Name	Field Description
Receiver signal average optical power	Displays the receiver signal average optical power, in milliwatts (mW) and decibels referred to 1.0 mW (dBm).
Laser bias current high alarm	Displays whether the laser bias power setting high alarm is **On** or **Off**.
Laser bias current low alarm	Displays whether the laser bias power setting low alarm is **On** or **Off**.
Laser bias current high warning	Displays whether the laser bias power setting high warning is **On** or **Off**.
Laser bias current low warning	Displays whether the laser bias power setting low warning is **On** or **Off**.
Laser output power high alarm	Displays whether the laser output power high alarm is **On** or **Off**.
Laser output power low alarm	Displays whether the laser output power low alarm is **On** or **Off**.
Laser output power high warning	Displays whether the laser output power high warning is **On** or **Off**.
Laser output power low warning	Displays whether the laser output power low warning is **On** or **Off**.
Module temperature high alarm	Displays whether the module temperature high alarm is **On** or **Off**.
Module temperature low alarm	Displays whether the module temperature low alarm is **On** or **Off**.
Module temperature high warning	Displays whether the module temperature high warning is **On** or **Off**.
Module temperature low warning	Displays whether the module temperature low warning is **On** or **Off**.
Module voltage high alarm	Displays whether the module voltage high alarm is **On** or **Off**.
Module voltage low alarm	Displays whether the module voltage low alarm is **On** or **Off**.
Module voltage high warning	Displays whether the module voltage high warning is **On** or **Off**.
Module voltage low warning	Displays whether the module voltage low warning is **On** or **Off**.
Laser rx power high alarm	Displays whether the receive laser power high alarm is **On** or **Off**.
Laser rx power low alarm	Displays whether the receive laser power low alarm is **On** or **Off**.
Laser rx power high warning	Displays whether the receive laser power high warning is **On** or **Off**.
Laser rx power low warning	Displays whether the receive laser power low warning is **On** or **Off**.
Laser bias current high alarm threshold	Displays the vendor-specified threshold for the laser bias current high alarm.
Laser bias current low alarm threshold	Displays the vendor-specified threshold for the laser bias current low alarm.

Table 136: show interfaces diagnostics optics Output Fields *(continued)*

Field Name	Field Description
Laser bias current high warning threshold	Displays the vendor-specified threshold for the laser bias current high warning.
Laser bias current low warning threshold	Displays the vendor-specified threshold for the laser bias current low warning.
Laser output power high alarm threshold	Displays the vendor-specified threshold for the laser output power high alarm.
Laser output power low alarm threshold	Displays the vendor-specified threshold for the laser output power low alarm.
Laser output power high warning threshold	Displays the vendor-specified threshold for the laser output power high warning.
Laser output power low warning threshold	Displays the vendor-specified threshold for the laser output power low warning.
Module temperature high alarm threshold	Displays the vendor-specified threshold for the module temperature high alarm.
Module temperature low alarm threshold	Displays the vendor-specified threshold for the module temperature low alarm.
Module temperature high warning threshold	Displays the vendor-specified threshold for the module temperature high warning.
Module temperature low warning threshold	Displays the vendor-specified threshold for the module temperature low warning.
Module voltage high alarm threshold	Displays the vendor-specified threshold for the module voltage high alarm.
Module voltage low alarm threshold	Displays the vendor-specified threshold for the module voltage low alarm.
Module voltage high warning threshold	Displays the vendor-specified threshold for the module voltage high warning.
Module voltage low warning threshold	Displays the vendor-specified threshold for the module voltage low warning.
Laser rx power high alarm threshold	Displays the vendor-specified threshold for the laser rx power high alarm.
Laser rx power low alarm threshold	Displays the vendor-specified threshold for the laser rx power low alarm.
Laser rx power high warning threshold	Displays the vendor-specified threshold for the laser rx power high warning.
Laser rx power low warning threshold	Displays the vendor-specified threshold for the laser rx power low warning.

show interfaces diagnostics optics ge-0/1/0 (SFP Transceiver)

```
user@host> show interfaces diagnostics optics ge-0/1/0
Physical interface: ge-0/1/0
    Laser bias current                        :  5.444 mA
    Laser output power                        :  0.3130 mW / -5.04 dBm
    Module temperature                        :  36 degrees C / 97 degrees F
    Module voltage                            :  3.2120 V
```

```
    Receiver signal average optical power     :  0.3840 mW / -4.16 dBm
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
    Module voltage high alarm                 :  Off
    Module voltage low alarm                  :  Off
    Module voltage high warning               :  Off
    Module voltage low warning                :  Off
    Laser rx power high alarm                 :  Off
    Laser rx power low alarm                  :  Off
    Laser rx power high warning               :  Off
    Laser rx power low warning                :  Off
    Laser bias current high alarm threshold   :  15.000 mA
    Laser bias current low alarm threshold    :  1.000 mA
    Laser bias current high warning threshold :  12.000 mA
    Laser bias current low warning threshold  :  2.000 mA
    Laser output power high alarm threshold   :  0.6300 mW / -2.01 dBm
    Laser output power low alarm threshold    :  0.0660 mW / -11.80 dBm
    Laser output power high warning threshold :  0.6300 mW / -2.01 dBm
    Laser output power low warning threshold  :  0.0780 mW / -11.08 dBm
    Module temperature high alarm threshold   :  109 degrees C / 228 degrees F
    Module temperature low alarm threshold    :  -29 degrees C / -20 degrees F
    Module temperature high warning threshold :  103 degrees C / 217 degrees F
    Module temperature low warning threshold  :  -13 degrees C / 9 degrees F
    Module voltage high alarm threshold       :  3.900 V
    Module voltage low alarm threshold        :  2.700 V
    Module voltage high warning threshold     :  3.700 V
    Module voltage low warning threshold      :  2.900 V
    Laser rx power high alarm threshold       :  1.2589 mW / 1.00 dBm
    Laser rx power low alarm threshold        :  0.0100 mW / -20.00 dBm
    Laser rx power high warning threshold     :  0.7939 mW / -1.00 dBm
    Laser rx power low warning threshold      :  0.0157 mW / -18.04 dBm
```

show interfaces diagnostics optics xe-0/1/0 (SFP+ Transceiver)

```
user@host> show interfaces diagnostics optics xe-0/1/0
Physical interface: xe-0/1/0
    Laser bias current                        :  4.968 mA
    Laser output power                        :  0.4940 mW / -3.06 dBm
    Module temperature                        :  27 degrees C / 81 degrees F
    Module voltage                            :  3.2310 V
    Receiver signal average optical power     :  0.0000
    Laser bias current high alarm             :  Off
    Laser bias current low alarm              :  Off
    Laser bias current high warning           :  Off
    Laser bias current low warning            :  Off
    Laser output power high alarm             :  Off
    Laser output power low alarm              :  Off
    Laser output power high warning           :  Off
    Laser output power low warning            :  Off
    Module temperature high alarm             :  Off
    Module temperature low alarm              :  Off
    Module temperature high warning           :  Off
    Module temperature low warning            :  Off
```

```
Module voltage high alarm                :  Off
Module voltage low alarm                 :  Off
Module voltage high warning              :  Off
Module voltage low warning               :  Off
Laser rx power high alarm                :  Off
Laser rx power low alarm                 :  On
Laser rx power high warning              :  Off
Laser rx power low warning               :  On
Laser bias current high alarm threshold  :  10.500 mA
Laser bias current low alarm threshold   :  2.000 mA
Laser bias current high warning threshold :  9.000 mA
Laser bias current low warning threshold  :  2.500 mA
Laser output power high alarm threshold  :  1.4120 mW / 1.50 dBm
Laser output power low alarm threshold   :  0.0740 mW / -11.31 dBm
Laser output power high warning threshold :  0.7070 mW / -1.51 dBm
Laser output power low warning threshold  :  0.1860 mW / -7.30 dBm
Module temperature high alarm threshold  :  75 degrees C / 167 degrees F
Module temperature low alarm threshold   :  -5 degrees C / 23 degrees F
Module temperature high warning threshold :  70 degrees C / 158 degrees F
Module temperature low warning threshold  :  0 degrees C / 32 degrees F
Module voltage high alarm threshold      :  3.630 V
Module voltage low alarm threshold       :  2.970 V
Module voltage high warning threshold    :  3.465 V
Module voltage low warning threshold     :  3.135 V
Laser rx power high alarm threshold      :  1.5849 mW / 2.00 dBm
Laser rx power low alarm threshold       :  0.0407 mW / -13.90 dBm
Laser rx power high warning threshold    :  0.7943 mW / -1.00 dBm
Laser rx power low warning threshold     :  0.1023 mW / -9.90 dBm
```

show interfaces ge-

Syntax

show interfaces ge-*fpc/pic/port*
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display status information about the specified Gigabit Ethernet interface.

Options

ge-*fpc/pic/port*—Display standard information about the specified Gigabit Ethernet interface.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index* —(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

Related Documentation

- Monitoring Interface Status and Traffic on page 931
- Troubleshooting Network Interfaces on J-EX4200 Switches on page 939
- Troubleshooting an Aggregated Ethernet Interface on page 940
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

List of Sample Output

show interfaces ge-0/0/0 on page 1012
show interfaces ge-0/0/0 brief on page 1012
show interfaces ge-0/0/0 detail on page 1012
show interfaces ge-0/0/4 extensive on page 1013

Output Fields

Table 137 on page 1005 lists the output fields for the **show interfaces ge-** command. Output fields are listed in the approximate order in which they appear.

Table 137: show interfaces ge- Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels

Table 137: show interfaces ge- Output Fields *(continued)*

Field Name	Field Description	Level of Output
Enabled	State of the interface: **Enabled** or **Disabled**.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Dell Support only (see "Requesting Technical Support" on page lxxi).	**detail extensive**
Description	Optional user-specified description.	**brief detail extensive**
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface. Default is 1514.	All levels
Speed	Speed at which the interface is running.	All levels
Loopback	Loopback status: **Enabled** or **Disabled**. If loopback is enabled, type of loopback: **Local** or **Remote**.	All levels
Source filtering	Source filtering status: **Enabled** or **Disabled**.	All levels
Flow control	Flow control status: **Enabled** or **Disabled**.	All levels
Auto-negotiation	Autonegotiation status: **Enabled** or **Disabled**.	All levels
Remote-fault	Remote fault status: • **Online**—Autonegotiation is manually configured as online. • **Offline**—Autonegotiation is manually configured as offline.	All levels
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels
Link flags	Information about the link.	All levels
CoS queues	Number of CoS queues configured.	**detail extensive** none
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
Current address	Configured MAC address.	**detail extensive** none
Hardware address	MAC address of the hardware.	**detail extensive** none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped:** ***year-month-day hour:minute:second timezone*** **(*hour:minute:second* ago)**. For example, **Last flapped: 2008–01–16 10:52:40 UTC (3d 22:58 ago).**	**detail extensive** none

Table 137: show interfaces ge- Output Fields *(continued)*

Field Name	Field Description	Level of Output
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface • **Output packets**—Number of packets transmitted on the interface. NOTE: The bandwidth bps counter is not enabled on the switch.	**detail extensive**
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • **L3 incompletes**—Number of incoming packets discarded because they failed Layer 3 sanity checks of the headers. For example, a frame with less than 20 bytes of available IP header is discarded. • **L2 channel errors**—Number of times the software did not find a valid logical interface for an incoming frame. • **L2 mismatch timeouts**—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • **FIFO errors**—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **Resource errors**—Sum of transmit drops.	**extensive**

Table 137: show interfaces ge- Output Fields *(continued)*

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • **Errors**—Sum of the outgoing frame aborts and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Collisions**—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • **Aged packets**—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • **FIFO errors**—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **HS link CRC errors**—Number of errors on the high-speed links between the ASICs responsible for handling the switch interfaces. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**
Egress queues	Total number of egress queues supported on the specified interface.	**detail extensive**
Queue counters (Egress)	CoS queue number and its associated user-configured forwarding class name. • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**
Active alarms and Active defects	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the switch configuration, an alarm can ring the red or yellow alarm bell on the switch or turn on the red or yellow alarm LED on the front of the switch. These fields can contain the value **None** or **Link**. • **None**—There are no active defects or alarms. • **Link**—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.	**detail extensive** none

Table 137: show interfaces ge- Output Fields *(continued)*

Field Name	Field Description	Level of Output
MAC statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC subsystem. • **Total octets** and **total packets**—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. • **Unicast packets, Broadcast packets,** and **Multicast packets**—Number of unicast, broadcast, and multicast packets. • **CRC/Align errors**—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • **FIFO error**—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **MAC control frames**—Number of MAC control frames. • **MAC pause frames**—Number of MAC control frames with **pause** operational code. • **Oversized frames**—Number of frames that exceed 1518 octets. • **Jabber frames**—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • **Fragment frames**—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • **Code violations**—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.”	**extensive**
Filter Statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC address filter subsystem.	**extensive**

Table 137: show interfaces ge- Output Fields *(continued)*

Field Name	Field Description	Level of Output

Autonegotiation information	Information about link autonegotiation:
	• **Negotiation status:**
	  • **Incomplete**—Ethernet interface has the speed or link mode configured.
	  • **No autonegotiation**—Remote Ethernet interface has the speed or link mode configured or does not perform autonegotiation.
	  • **Complete**—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• **Link partner status**—**OK** when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• **Link partner:**
	  • **Link mode**—Depending on the capability of the attached Ethernet device, either **Full-duplex** or **Half-duplex**.
	  • **Flow control**—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are **Symmetric** (link partner supports **PAUSE** on receive and transmit), **Asymmetric** (link partner supports **PAUSE** on transmit), and **Symmetric/Asymmetric** (link partner supports **PAUSE** on both receive and transmit or **PAUSE** only on receive).
	  • **Remote fault**—Remote fault information from the link partner—**Failure** indicates a receive link error. **OK** indicates that the link partner is receiving. **Negotiation error** indicates a negotiation error. **Offline** indicates that the link partner is going offline.
	  • **Link partner speed**—Speed of the link partner.
	• **Local resolution**—Information from the link partner:
	  • **Flow control**—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are **Symmetric** (link partner supports **PAUSE** on receive and transmit), **Asymmetric** (link partner supports **PAUSE** on transmit), and **Symmetric/Asymmetric** (link partner supports **PAUSE** on both receive and transmit or **PAUSE** only on receive).
	  • **Remote fault**—Remote fault information. **Link OK** (no error detected on receive), **Offline** (local interface is offline), and **Link Failure** (link error detected on receive).
		extensive

Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • **Destination slot**—FPC slot number. NOTE: For a J-EX4200 standalone switch, the FPC slot number refers to the switch itself and is always **0**. In a Virtual Chassis configuration, the FPC slot number refers to the member ID. In a J-EX8200 switch, the FPC slot number refers to the line card slot number.	**extensive**
Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP interface index number for the logical interface.	**detail extensive** none

Table 137: show interfaces ge- Output Fields *(continued)*

Field Name	Field Description	Level of Output
Generation	Unique number for use by Dell Support only (see "Requesting Technical Support" on page lxxi).	**detail extensive**
Flags	Information about the logical interface.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family.	**detail extensive** none
Traffic statistics	Number and rate of bytes and packets received (input) and transmitted (output) on the specified interface.	**detail extensive**
IPv6 transit statistics	If IPv6 statistics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface.	**extensive**
Local statistics	Number and rate of bytes and packets destined to and from the switch.	**extensive**
Transit statistics	Number and rate of bytes and packets transiting the switch.	**extensive**
Generation	Unique number for use by Dell Support only (see "Requesting Technical Support" on page lxxi).	**detail extensive**
Route Table	Route table in which the logical interface address is located. For example, **0** refers to the routing table **inet.0**.	**detail extensive** none
Input Filters	Names of any input filters applied to this interface.	**detail extensive**
Output Filters	Names of any output filters applied to this interface.	**detail extensive**
Flags	Information about protocol family flags. If unicast reverse-path forwarding (RPF) is explicitly configured on the specified interface, the uRPF flag is displayed. If unicast RPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag is not displayed even though unicast RPF is enabled.	**detail extensive**
protocol-family	Protocol family configured on the logical interface. If the protocol is **inet**, the IP address of the interface is also displayed.	**brief**
Flags	Information about the address flags.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address of the logical interlace.	**detail extensive** none

Table 137: show interfaces ge- Output Fields *(continued)*

Field Name	Field Description	Level of Output
Generation	Unique number for use by Dell Support only (see "Requesting Technical Support" on page lxxi).	**detail extensive**

show interfaces ge-0/0/0

```
user@switch> show interfaces ge-0/0/0
Physical interface: ge-0/0/0, Enabled, Physical link is Down
  Interface index: 129, SNMP ifIndex: 21
  Link-level type: Ethernet, MTU: 1514, Speed: Unspecified, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled
  Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:19:e2:50:3f:41, Hardware address: 00:19:e2:50:3f:41
  Last flapped   : 2008-01-16 11:40:53 UTC (4d 02:30 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Ingress rate at Packet Forwarding Engine      : 0 bps (0 pps)
  Ingress drop rate at Packet Forwarding Engine : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None

  Logical interface ge-0/0/0.0 (Index 65) (SNMP ifIndex 22)
    Flags: SNMP-Traps
    Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Protocol eth-switch
      Flags: None
```

show interfaces ge-0/0/0 brief

```
user@switch> show interfaces ge-0/0/0 brief
Physical interface: ge-0/0/0, Enabled, Physical link is Down
  Description: voice priority and tcp and icmp traffic rate-limiting filter at i
ngress port
  Link-level type: Ethernet, MTU: 1514, Speed: Unspecified, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running Down
  Interface flags: Hardware-Down SNMP-Traps Internal: 0x0
  Link flags     : None

  Logical interface ge-0/0/0.0
    Flags: Device-Down SNMP-Traps Encapsulation: ENET2
    eth-switch
```

show interfaces ge-0/0/0 detail

```
user@switch> show interfaces ge-0/0/0 detail
Physical interface: ge-0/0/0, Enabled, Physical link is Up
  Interface index: 193, SNMP ifIndex: 206, Generation: 196
  Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Auto,
  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
  Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
```

```
CoS queues     : 8 supported, 8 maximum usable queues
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:1f:12:30:ff:40, Hardware address: 00:1f:12:30:ff:40
Last flapped   : 2009-05-05 06:03:05 UTC (00:22:13 ago)
Statistics last cleared: Never
Traffic statistics:
 Input  bytes  :                    0                    0 bps
 Output bytes  :                    0                    0 bps
 Input  packets:                    0                    0 pps
 Output packets:                    0                    0 pps
 IPv6 transit statistics:
  Input  bytes  :                   0
  Output bytes  :                   0
  Input  packets:                   0
  Output packets:                   0
Egress queues: 8 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets

  0 best-effort                   0                    0                    0

  1 assured-forw                  0                    0                    0

  5 expedited-fo                  0                    0                    0

  7 network-cont                  0                    0                    0

Active alarms  : None
Active defects : None

Logical interface ge-0/0/0.0 (Index 65) (SNMP ifIndex 235) (Generation 130)
  Flags: SNMP-Traps Encapsulation: ENET2
  Bandwidth: 0
  Traffic statistics:
   Input  bytes  :                   0
   Output bytes  :                   0
   Input  packets:                   0
   Output packets:                   0
  Local statistics:
   Input  bytes  :                   0
   Output bytes  :                   0
   Input  packets:                   0
   Output packets:                   0
  Transit statistics:
   Input  bytes  :                   0                    0 bps
   Output bytes  :                   0                    0 bps
   Input  packets:                   0                    0 pps
   Output packets:                   0                    0 pps
  Protocol eth-switch, Generation: 146, Route table: 0
    Flags: Is-Primary
    Input Filters: f1,
    Output Filters: f2,,,,
```

show interfaces ge-0/0/4 extensive

```
user@switch> show interfaces ge-0/0/4 extensive
Physical interface: ge-0/0/4, Enabled, Physical link is Up
  Interface index: 165, SNMP ifIndex: 152, Generation: 168
  Link-level type: Ethernet, MTU: 1514, Speed: Auto, Duplex: Auto,
  MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled,
  Flow control: Enabled, Auto-negotiation: Enabled, Remote fault: Online
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
```

```
CoS queues     : 8 supported, 8 maximum usable queues
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:1f:12:33:65:44, Hardware address: 00:1f:12:33:65:44
Last flapped   : 2008-09-17 11:02:25 UTC (16:32:54 ago)
Statistics last cleared: Never
Traffic statistics:
 Input  bytes  :                    0                    0 bps
 Output bytes  :              2989761                  984 bps
 Input  packets:                    0                    0 pps
 Output packets:                24307                    1 pps
 IPv6 transit statistics:
  Input  bytes  :                   0
  Output bytes  :                   0
  Input  packets:                   0
  Output packets:                   0
Input errors:
  Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
  L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
  FIFO errors: 0, Resource errors: 0
Output errors:
  Carrier transitions: 1, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

  FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Egress queues: 8 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets

  0 best-effort                    0                    0                    0

  1 assured-forw                   0                    0                    0

  5 expedited-fo                   0                    0                    0

  7 network-cont                   0                24307                    0

Active alarms  : None
Active defects : None
MAC statistics:                       Receive         Transmit
  Total octets                              0          2989761
  Total packets                             0            24307
  Unicast packets                           0                0
  Broadcast packets                         0                0
  Multicast packets                         0            24307
  CRC/Align errors                          0                0
  FIFO errors                               0                0
  MAC control frames                        0                0
  MAC pause frames                          0                0
  Oversized frames                          0
  Jabber frames                             0
  Fragment frames                           0
  Code violations                           0
Autonegotiation information:
  Negotiation status: Complete
  Link partner:
      Link mode: Full-duplex, Flow control: None, Remote fault: OK,
      Link partner Speed: 1000 Mbps
  Local resolution:
      Flow control: None, Remote fault: Link OK
Packet Forwarding Engine configuration:
  Destination slot: 0
  Direction : Output
  CoS transmit queue               Bandwidth               Buffer Priority
```

```
Limit
                              %           bps      %           usec
    0 best-effort            95     950000000     95             NA       low
none
    7 network-control         5      50000000      5             NA       low
none

  Logical interface ge-0/0/4.0 (Index 82) (SNMP ifIndex 184) (Generation 147)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                   0
     Output bytes  :             4107883
     Input  packets:                   0
     Output packets:               24307
     IPv6 transit statistics:
      Input  bytes  :                  0
      Output bytes  :                  0
      Input  packets:                  0
      Output packets:                  0
    Local statistics:
     Input  bytes  :                   0
     Output bytes  :             4107883
     Input  packets:                   0
     Output packets:               24307
    Transit statistics:
     Input  bytes  :                   0                  0 bps
     Output bytes  :                   0                  0 bps
     Input  packets:                   0                  0 pps
     Output packets:                   0                  0 pps
     IPv6 transit statistics:
      Input  bytes  :                  0
      Output bytes  :                  0
      Input  packets:                  0
      Output packets:                  0
    Protocol eth-switch, Generation: 159, Route table: 0
      Flags: None
      Input Filters: f2,
      Output Filters: f1,,,,
```

show interfaces queue

Syntax

show interfaces queue
<both-ingress-egress>
<egress>
<forwarding-class *forwarding-class*>
<ingress>
<interface-name *interface-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display class-of-service (CoS) queue information for physical interfaces.

Options

none—Show detailed CoS queue statistics for all physical interfaces.

both-ingress-egress—(Optional) Show both ingress and egress queue statistics. (Ingress statistics are not available for all interfaces.)

egress—(Optional) Show eqress queue statistics only.

forwarding-class *forwarding-class*—(Optional) Show queue statistics only for the specified forwarding class.

ingress—(Optional) Show ingress queue statistics only. (Ingress statistics are not available for all interfaces.)

interface-name *interface-name*—(Optional) Show queue statistics for the specified interface.

Required Privilege Level

view

Related Documentation

- Monitoring Interface Status and Traffic on page 931
- Monitoring Interfaces That Have CoS Components on page 2937
- Defining CoS Schedulers (CLI Procedure) on page 2920
- Configuring CoS Traffic Classification for Ingress Queuing on 40-port SFP+ Line Cards (CLI Procedure)

List of Sample Output

show interfaces queue ge-0/0/0 on page 1018

Output Fields

Table 138 on page 1016 lists the output fields for the **show interfaces queue** command. Output fields are listed in the approximate order in which they appear.

Table 138: show interfaces queue Output Fields

Field Name	Field Description
Physical Interface and Forwarding Class Information	
Physical interface	Name of the physical interface.

Table 138: show interfaces queue Output Fields *(continued)*

Field Name	Field Description
Enabled	State of the interface. Possible values are: • **Administratively down, Physical link is Down**—The interface is turned off, and the physical link is inoperable. • **Administratively down, Physical link is Up**—The interface is turned off, but the physical link is operational and can pass packets when it is enabled. • **Enabled, Physical link is Down**—The interface is turned on, but the physical link is inoperable and cannot pass packets. • **Enabled, Physical link is Up**—The interface is turned on, and the physical link is operational and can pass packets.
Interface index	Index number of the physical interface, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the physical interface.
Description	User-configured interface description.
Forwarding classes	Number of forwarding classes supported and in use for the interface.
Ingress Queues Information (not shown for all interfaces)	
Ingress queues	Number of input queues supported and in use on the specified interface.
Transmitted	Transmission statistics for the queue: • **Packets**—Number of packets transmitted by this queue. • **Bytes**—Number of bytes transmitted by this queue. • **Tail-dropped packets**—Number of packets dropped because the queue buffers were full.
PFE chassis queues	For an interface on an oversubscribed line card, the number of Packet Forwarding Engine chassis queues supported and in use for the port group to which the interface belongs. The Packet Forwarding Engine chassis queue for a port group handles high priority traffic from all the interfaces in the port group.
Egress Queues Information	
Egress queues	Number of output queues supported and in use on the specified interface.
Queue	CoS queue number.
Queued	This counter is not supported on J-EX Series switches.

Table 138: show interfaces queue Output Fields *(continued)*

Field Name	Field Description

Transmitted	Number of packets and bytes transmitted by this queue. Information on transmitted packets and bytes can include:
	• **Packets**—Number of packets transmitted.
	• **Bytes**—Number of bytes transmitted.
	• **Tail-dropped packets**—Number of arriving packets dropped because output queue buffers were full.
	• **RED-dropped packets**—Number of packets dropped because of random early detection (RED).
	  ◦ **Low**—Number of low loss priority packets dropped because of RED.
	  ◦ **High**—Number of high loss priority packets dropped because of RED.
	• **RED-dropped bytes**—Number of bytes dropped because of random early detection (RED).
	  ◦ **Low**—Number of low loss priority bytes dropped because of RED.
	  ◦ **High**—Number of high loss priority bytes dropped because of RED.

Packet Forwarding Engine Chassis Queues	For an interface on an oversubscribed line card, the number of Packet Forwarding Engine chassis queues supported and in use for the port group to which the interface belongs. The queue statistics reflect the traffic flowing on all the interfaces in the port group.

show interfaces queue ge-0/0/0

```
user@switch> show interfaces queue ge-0/0/0
Physical interface: ge-0/0/0, Enabled, Physical link is Down
  Interface index: 130, SNMP ifIndex: 501
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
  Transmitted:
    Packets              :                    0
    Bytes                :                    0
    Tail-dropped packets :                    0
Queue: 1, Forwarding classes: assured-forwarding
  Queued:
  Transmitted:
    Packets              :                    0
    Bytes                :                    0
    Tail-dropped packets :                    0
Queue: 5, Forwarding classes: expedited-forwarding
  Queued:
  Transmitted:
    Packets              :                    0
    Bytes                :                    0
    Tail-dropped packets :                    0
Queue: 7, Forwarding classes: network-control
  Queued:
  Transmitted:
    Packets              :                    0
    Bytes                :                    0
    Tail-dropped packets :                    0
```

show interfaces xe-

Syntax

show interfaces xe-*fpc/pic/port*
<brief | detail | extensive | terse>
<descriptions>
<media>
<snmp-index *snmp-index*>
<statistics>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display status information about the specified 10-Gigabit Ethernet interface.

Options

xe-*fpc/pic/port* —Display standard information about the specified 10-Gigabit Ethernet interface.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

descriptions—(Optional) Display interface description strings.

media—(Optional) Display media-specific information about network interfaces.

snmp-index *snmp-index* —(Optional) Display information for the specified SNMP index of the interface.

statistics—(Optional) Display static interface statistics.

Required Privilege Level

view

Related Documentation

- Monitoring Interface Status and Traffic on page 931
- Troubleshooting Network Interfaces on J-EX4200 Switches on page 939
- Troubleshooting an Aggregated Ethernet Interface on page 940
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

List of Sample Output

show interfaces xe-0/1/0 on page 1026
show interfaces xe-4/1/0 on page 1027
show interfaces xe-0/1/0 brief on page 1027
show interfaces xe-4/1/0 detail on page 1027
show interfaces xe-4/1/0 extensive on page 1028

Output Fields

Table 139 on page 1020 lists the output fields for the **show interfaces xe-** command. Output fields are listed in the approximate order in which they appear.

Table 139: show interfaces xe- Output Fields

Field Name	Field Description	Level of Output
Physical Interface		
Physical interface	Name of the physical interface.	All levels
Enabled	State of the interface.	All levels
Interface index	Index number of the physical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP index number for the physical interface.	**detail extensive** none
Generation	Unique number for use by Dell Support only (see "Requesting Technical Support" on page lxxi).	**detail extensive**
Link-level type	Encapsulation being used on the physical interface.	All levels
MTU	Maximum transmission unit size on the physical interface.	All levels
Speed	Speed at which the interface is running.	All levels
Duplex	Duplicity of the interface.	All levels
BPDU Error	blah blah	All levels
Loopback	Loopback status: **Enabled** or **Disabled**. If loopback is enabled, type of loopback: **Local** or **Remote**.	All levels
Source filtering	Source filtering status: **Enabled** or **Disabled**.	All levels
Flow control	Flow control status: **Enabled** or **Disabled**.	All levels
Device flags	Information about the physical device.	All levels
Interface flags	Information about the interface.	All levels
Link flags	Information about the link.	All levels
Wavelength	Configured wavelength, in nanometers (nm).	All levels
Frequency	Frequency associated with the configured wavelength, in terahertz (THz).	All levels
CoS queues	Number of CoS queues configured.	**detail extensive** none
Schedulers	Number of CoS schedulers configured.	**extensive**
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.	**detail extensive**
Current address	Configured MAC address.	**detail extensive** none

Table 139: show interfaces xe- Output Fields *(continued)*

Field Name	Field Description	Level of Output
Hardware address	Hardware MAC address.	**detail extensive** none
Last flapped	Date, time, and how long ago the interface went from down to up. The format is **Last flapped: *year-month-day hour::minute*:*second*: *timezone* (*day*d *hour:minute*:*second* ago)**. For example, **Last flapped: 2008–01–16 10:52:40 UTC (3d 22:58 ago).**	**detail extensive** none
Input Rate	Input rate in bits per second (bps) and packets per second (pps).	None specified
Output Rate	Output rate in bps and pps.	None specified
Statistics last cleared	Time when the statistics for the interface were last set to zero.	**detail extensive**
Traffic statistics	Number and rate of bytes and packets received and transmitted on the physical interface. • **Input bytes**—Number of bytes received on the interface. • **Output bytes**—Number of bytes transmitted on the interface. • **Input packets**—Number of packets received on the interface • **Output packets**—Number of packets transmitted on the interface. NOTE: The bandwidth bps counter is not enabled on the switch.	**detail extensive**
Input errors	Input errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Errors**—Sum of the incoming frame aborts and FCS errors. • **Drops**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Framing errors**—Number of packets received with an invalid frame checksum (FCS). • **Runts**—Number of frames received that are smaller than the runt threshold. • **Policed discards**—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that the Junos OS does not handle. • **L3 incompletes**—Number of incoming packets discarded because they failed Layer 3 sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. L3 incomplete errors can be ignored if you configure the **ignore-l3-incompletes** statement. • **L2 channel errors**—Number of times the software did not find a valid logical interface for an incoming frame. • **L2 mismatch timeouts**—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • **FIFO errors**—Number of FIFO errors in the receive direction that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **Resource errors**—Sum of transmit drops.	**extensive**

Table 139: show interfaces xe- Output Fields *(continued)*

Field Name	Field Description	Level of Output
Output errors	Output errors on the interface. The following paragraphs explain the counters whose meaning might not be obvious: • **Carrier transitions**—Number of times the interface has gone from **down** to **up**. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and then up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC or PIM is malfunctioning. • **Errors**—Sum of the outgoing frame aborts and FCS errors. • **Drops**—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Collisions**—Number of Ethernet collisions. The Gigabit Ethernet PIC supports only full-duplex operation, so for Gigabit Ethernet PICs, this number should always remain 0. If it is nonzero, there is a software bug. • **Aged packets**—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • **FIFO errors**—Number of FIFO errors in the send direction as reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **HS link CRC errors**—Number of errors on the high-speed links between the ASICs responsible for handling the switch interfaces. • **MTU errors**—Number of packets whose size exceeded the MTU of the interface. • **Resource errors**—Sum of transmit drops.	**extensive**
Ingress queues	Total number of ingress queues supported on the specified interface.	**extensive**
Queue counters (Ingress)	Statistics for the CoS low and high priority ingress queues: • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**extensive**
Egress queues	Total number of egress queues supported on the specified interface.	**detail extensive**
Queue counters (Egress)	Statistics for the CoS egress queues: • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism.	**detail extensive**
PFE Egress queues	Total number of Packet Forwarding Engine egress queues shared by the interfaces in a port group. Displayed for an interface on a 40–port SFP+ line card.	**detail extensive**

Table 139: show interfaces xe- Output Fields *(continued)*

Field Name	Field Description	Level of Output
Queue counters (Packet Forwarding Engine Egress)	Statistics for the Packet Forwarding Engine egress queues: • **Queued packets**—Number of queued packets. • **Transmitted packets**—Number of transmitted packets. • **Dropped packets**—Number of packets dropped by the ASIC's RED mechanism. Displayed for an interface on a 40–port SFP+ line card.	**detail extensive**
Active alarms and Active defects	Ethernet-specific defects that can prevent the interface from passing packets. When a defect persists for a certain amount of time, it is promoted to an alarm. Based on the switch configuration, an alarm can ring the red or yellow alarm bell on the switch or turn on the red or yellow alarm LED on the front of the switch. These fields can contain the value **None** or **Link**. • **None**—There are no active defects or alarms. • **Link**—Interface has lost its link state, which usually means that the cable is unplugged, the far-end system has been turned off, or the PIC is malfunctioning.	**detail extensive** none
PCS statistics	Physical Coding Sublayer (PCS) fault conditions from the LAN PHY device.	**detail extensive**
MAC statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC subsystem. • **Total octets** and **total packets**—Total number of octets and packets. For Gigabit Ethernet IQ PICs, the received octets count varies by interface type. • **Unicast packets, Broadcast packets,** and **Multicast packets**—Number of unicast, broadcast, and multicast packets. • **CRC/Align errors**—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS Error) or a bad FCS with a nonintegral number of octets (Alignment Error). • **FIFO error**—Number of FIFO errors that are reported by the ASIC on the PIC. If this value is ever nonzero, the PIC is probably malfunctioning. • **MAC control frames**—Number of MAC control frames. • **MAC pause frames**—Number of MAC control frames with **pause** operational code. • **Oversized frames**—Number of frames that exceed 1518 octets. • **Jabber frames**—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • **Fragment frames**—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • **Code violations**—Number of times an event caused the PHY to indicate “Data reception error” or “invalid data symbol error.”	**extensive**

Table 139: show interfaces xe- Output Fields *(continued)*

Field Name	Field Description	Level of Output
Filter statistics	**Receive** and **Transmit** statistics reported by the PIC's MAC address filter subsystem.	**extensive**

Autonegotiation information	Information about link autonegotiation:
	• **Negotiation status:**
	  ◦ **Incomplete**—Ethernet interface has the speed or link mode configured.
	  ◦ **No autonegotiation**—Remote Ethernet interface has the speed or link mode configured or does not perform autonegotiation.
	  ◦ **Complete**—Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• **Link partner status**—**OK** when Ethernet interface is connected to a device that performs autonegotiation and the autonegotiation process is successful.
	• **Link partner:**
	  ◦ **Link mode**—Depending on the capability of the attached Ethernet device, either **Full-duplex** or **Half-duplex**.
	  ◦ **Flow control**—Types of flow control supported by the remote Ethernet device. For Fast Ethernet interfaces, the type is **None**. For Gigabit Ethernet interfaces, types are **Symmetric** (link partner supports **PAUSE** on receive and transmit), **Asymmetric** (link partner supports **PAUSE** on transmit), and **Symmetric/Asymmetric** (link partner supports **PAUSE** on both receive and transmit or **PAUSE** only on receive).
	  ◦ **Remote fault**—Remote fault information from the link partner—**Failure** indicates a receive link error. **OK** indicates that the link partner is receiving. **Negotiation error** indicates a negotiation error. **Offline** indicates that the link partner is going offline.
	• **Local resolution**—Information from the link partner:
	  ◦ **Flow control**—Types of flow control supported by the remote Ethernet device. For Gigabit Ethernet interfaces, types are **Symmetric** (link partner supports **PAUSE** on receive and transmit), **Asymmetric** (link partner supports **PAUSE** on transmit), and **Symmetric/Asymmetric** (link partner supports **PAUSE** on both receive and transmit or **PAUSE** only on receive).
	  ◦ **Remote fault**—Remote fault information. **Link OK** (no error detected on receive), **Offline** (local interface is offline), and **Link Failure** (link error detected on receive).
		extensive

Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: • **Destination slot**—FPC slot number. NOTE: For a J-EX4200 standalone switch, the FPC slot number refers to the switch itself and is always **0**. In a Virtual Chassis configuration, FPC slot number refers to the member ID. In a J-EX8200 switch, the FPC slot number refers to the line card slot number.	**extensive**

Table 139: show interfaces xe- Output Fields *(continued)*

Field Name	Field Description	Level of Output

CoS Information	Information about the CoS queue for the physical interface:
	• **Direction**—Queue direction, either **Input** or **Output**.
	• **CoS transmit queue**—Queue number and its associated user-configured forwarding class name.
	• **Bandwidth**—Information about bandwidth allocated to the queue:
	  • **%**—Bandwidth allocated to the queue as a percentage
	  • **bps**—Bandwidth allocated to the queue in bps
	• **Buffer**—Information about buffer space allocated to the queue:
	  • **%**—Buffer space allocated to the queue as a percentage.
	  • **usec**—Buffer space allocated to the queue in microseconds. This value is nonzero only if the buffer size is configured in terms of time.
	• **Priority**—Queue priority: **low** or **high**.
	• **Limit**—Displayed if rate limiting is configured for the queue. Possible values are **none** and **exact**. If **exact** is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If **none** is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.	

Logical Interface		
Logical interface	Name of the logical interface.	All levels
Index	Index number of the logical interface, which reflects its initialization sequence.	**detail extensive** none
SNMP ifIndex	SNMP interface index number for the logical interface.	**detail extensive** none
Generation	Unique number for use by Dell technical support only.	**detail extensive**
Flags	Information about the logical interface.	All levels
Encapsulation	Encapsulation on the logical interface.	All levels
Protocol	Protocol family.	**detail extensive** none
Traffic statistics	Number and rate of bytes and packets received (input) and transmitted (output) on the specified interface.	**detail extensive**
IPv6 transit statistics	If IPv6 statistics tracking is enabled, number of IPv6 bytes and packets received and transmitted on the logical interface.	**extensive**
Local statistics	Number and rate of bytes and packets destined to and from the switch.	**extensive**
Transit statistics	Number and rate of bytes and packets transiting the switch.	**extensive**
Generation	Unique number for use by Dell Support only (see "Requesting Technical Support" on page lxxi).	**detail extensive**

Table 139: show interfaces xe- Output Fields *(continued)*

Field Name	Field Description	Level of Output
Route Table	Route table in which the logical interface address is located. For example, **0** refers to the routing table **inet.0**.	**detail extensive** none
Input Filters	Names of any input filters applied to this interface.	**detail extensive**
Output Filters	Names of any output filters applied to this interface.	**detail extensive**
Flags	Information about protocol family flags. If unicast reverse-path forwarding (RPF) is explicitly configured on the specified interface, the uRPF flag is display. If unicast RPF was configured on a different interface (and therefore is enabled on all switch interfaces) but was not explicitly configured on the specified interface, the uRPF flag is not displayed even though unicast RPF is enabled.	**detail extensive**
Addresses, Flags	Information about the address flags.	**detail extensive** none
protocol-family	Protocol family configured on the logical interface. If the protocol is **inet**, the IP address of the interface is also displayed.	**brief**
Flags	Information about the address flags.	**detail extensive** none
Destination	IP address of the remote side of the connection.	**detail extensive** none
Local	IP address of the logical interface.	**detail extensive** none
Broadcast	Broadcast address of the logical interlace.	**detail extensive** none
Generation	Unique number for use by Dell Support only (see "Requesting Technical Support" on page lxxi).	**detail extensive**

show interfaces xe-0/1/0

```
user@switch> show interfaces xe-0/1/0
Physical interface: xe-0/1/0, Enabled, Physical link is Up
  Interface index: 153, SNMP ifIndex: 69
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:19:e2:50:c8:99, Hardware address: 00:19:e2:50:c8:99
  Last flapped   : 2008-02-25 05:28:08 UTC (00:12:49 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None
  Logical interface xe-0/1/0.0 (Index 88) (SNMP ifIndex 70)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Protocol eth-switch
```

```
    Flags: None
```

show interfaces xe-4/1/0

```
user@switch show interfaces xe-4/1/0
Physical interface: xe-4/1/0, Enabled, Physical link is Up
  Interface index: 387, SNMP ifIndex: 369
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex,
  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Current address: 00:23:9c:03:8e:70, Hardware address: 00:23:9c:03:8e:70
  Last flapped   : 2009-05-12 08:01:04 UTC (00:13:44 ago)
  Input rate     : 36432 bps (3 pps)
  Output rate    : 0 bps (0 pps)
  Active alarms  : None
  Active defects : None

  Logical interface xe-4/1/0.0 (Index 66) (SNMP ifIndex 417)
    Flags: SNMP-Traps Encapsulation: ENET2
    Input packets : 0
    Output packets: 0
    Protocol eth-switch
      Flags: None
```

show interfaces xe-0/1/0 brief

```
user@switch> show interfaces xe-0/1/0 brief
Physical interface: xe-0/1/0, Enabled, Physical link is Up
  Link-level type: Ethernet, MTU: 1514, Speed: 1000mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None

  Logical interface xe-0/1/0.0
    Flags: SNMP-Traps Encapsulation: ENET2
    eth-switch
```

show interfaces xe-4/1/0 detail

```
user@switch> show interfaces xe-4/1/0 detail
Physical interface: xe-4/1/0, Enabled, Physical link is Up
  Interface index: 387, SNMP ifIndex: 369, Generation: 390
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex,
  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:23:9c:03:8e:70, Hardware address: 00:23:9c:03:8e:70
  Last flapped   : 2009-05-12 08:01:04 UTC (00:13:49 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              4945644                48576 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                 3258                    4 pps
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
```

```
  Output bytes  :                   0
  Input  packets:                   0
  Output packets:                   0
Egress queues: 8 supported, 4 in use
Queue counters:       Queued packets  Transmitted packets      Dropped packets

  0 best-effort                    0                    0                    0

  1 assured-forw                   0                    0                    0

  5 expedited-fo                   0                    0                    0

  7 network-cont                   0                    0                    0

Active alarms  : None
Active defects : None

Logical interface xe-4/1/0.0 (Index 66) (SNMP ifIndex 417) (Generation 158)
  Flags: SNMP-Traps Encapsulation: ENET2
  Traffic statistics:
   Input  bytes  :                   0
   Output bytes  :                   0
   Input  packets:                   0
   Output packets:                   0
  Local statistics:
   Input  bytes  :                   0
   Output bytes  :                   0
   Input  packets:                   0
   Output packets:                   0
  Transit statistics:
   Input  bytes  :                   0                   0 bps
   Output bytes  :                   0                   0 bps
   Input  packets:                   0                   0 pps
   Output packets:                   0                   0 pps
  Protocol eth-switch, Generation: 174, Route table: 0
    Flags: None
    Input Filters: f1,
    Output Filters: f2,,,,
```

show interfaces xe-4/1/0 extensive

```
user@switch> show interfaces xe-4/1/0 extensive
Physical interface: xe-4/1/0, Enabled, Physical link is Up
  Interface index: 387, SNMP ifIndex: 369, Generation: 390
  Link-level type: Ethernet, MTU: 1514, Speed: 10Gbps, Duplex: Full-Duplex,
  BPDU Error: None, MAC-REWRITE Error: None, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x0
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:23:9c:03:8e:70, Hardware address: 00:23:9c:03:8e:70
  Last flapped   : 2009-05-12 08:01:04 UTC (00:14:01 ago)
  Statistics last cleared: Never
  Traffic statistics:
   Input  bytes  :              5015472                36432 bps
   Output bytes  :                    0                    0 bps
   Input  packets:                 3304                    3 pps
   Output packets:                    0                    0 pps
   IPv6 transit statistics:
    Input  bytes  :                   0
    Output bytes  :                   0
```

```
    Input  packets:                  0
    Output packets:                  0
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0,
    L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0,
    FIFO errors: 0, Resource errors: 0
  Output errors:
    Carrier transitions: 3, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,

    FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
  Egress queues: 8 supported, 4 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets

    0 best-effort                   0                    0                    0

    1 assured-forw                  0                    0                    0

    5 expedited-fo                  0                    0                    0

    7 network-cont                  0                    0                    0

  Active alarms  : None
  Active defects : None
  MAC statistics:                      Receive         Transmit
    Total octets                       5015472                0
    Total packets                         3304                0
    Unicast packets                       3304                0
    Broadcast packets                        0                0
    Multicast packets                        0                0
    CRC/Align errors                         0                0
    FIFO errors                              0                0
    MAC control frames                       0                0
    MAC pause frames                         0                0
    Oversized frames                         0
    Jabber frames                            0
    Fragment frames                          0
    Code violations                          0
  Packet Forwarding Engine configuration:
    Destination slot: 4
    Direction : Output
    CoS transmit queue               Bandwidth               Buffer Priority
Limit
                               %            bps     %           usec
    0 best-effort             95     9500000000    95             NA      low
none
    7 network-control          5      500000000     5             NA      low
none

  Logical interface xe-4/1/0.0 (Index 66) (SNMP ifIndex 417) (Generation 158)
    Flags: SNMP-Traps Encapsulation: ENET2
    Traffic statistics:
     Input  bytes  :                  0
     Output bytes  :                  0
     Input  packets:                  0
     Output packets:                  0
    Local statistics:
     Input  bytes  :                  0
     Output bytes  :                  0
     Input  packets:                  0
     Output packets:                  0
    Transit statistics:
```

```
 Input  bytes  :                    0                    0 bps
 Output bytes  :                    0                    0 bps
 Input  packets:                    0                    0 pps
 Output packets:                    0                    0 pps
Protocol eth-switch, Generation: 174, Route table: 0
  Flags: None
  Input Filters: f1,
  Output Filters: f2,,,,
```

show ipv6 neighbors

Syntax show ipv6 neighbors

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display information about the IPv6 neighbor cache.

Options This command has no options.

Required Privilege Level view

Related Documentation
- **clear ipv6 neighbors on page 990**

List of Sample Output
show ipv6 neighbors on page 1031
show ipv6 neighbors on page 1031

Output Fields Table 140 on page 1031 describes the output fields for the **show ipv6 neighbors** command. Output fields are listed in the approximate order in which they appear.

Table 140: show ipv6 neighbors Output Fields

Field Name	Field Description
IPv6 Address	Name of the IPv6 interface.
Linklayer Address	Link-layer address.
State	State of the link: **up**, **down**, **incomplete**, **reachable**, **stale**, or **unreachable**.
Exp	Number of seconds until the entry expires.
Rtr	Whether the neighbor is a routing device: **yes** or **no**.
Secure	Whether this entry was created using the Secure Neighbor Discovery (SEND) protocol: **yes** or **no**.
Interface	Name of the interface.

show ipv6 neighbors

```
user@host> show ipv6 neighbors
IPv6 Address                 Linklayer Address  State       Exp  Rtr  Interface
fe80::2a0:c9ff:fe5b:4c1e     00:a0:c9:5b:4c:1e  reachable   15   yes  fxp0.0
```

show ipv6 neighbors

```
user@host > show ipv6 neighbors
IPv6 Address                 Linklayer Address  State       Exp Rtr Secure
Interface
```

```
fe80::14fb:5dcf:54bd:ff76    00:90:69:a0:a8:bc  stale       1113 yes yes
ge-3/2/0.0
```

show lacp interfaces

Syntax show lacp interfaces *interface-name*

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display Link Aggregation Control Protocol (LACP) information about the specified aggregated Ethernet or Gigabit Ethernet interface.

Options none—Display LACP information for all interfaces.

interface-name—(Optional) Display LACP information for the specified interface:

- Aggregated Ethernet—**ae*x***
- Gigabit Ethernet—**ge-*fpc/pic/port***

Required Privilege Level view

Related Documentation

- Example: Configuring Aggregated Ethernet High-Speed Uplinks Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 740
- Example: Configuring Aggregated Ethernet High-Speed Uplinks with LACP Between a Virtual Chassis Access Switch and a Virtual Chassis Distribution Switch on page 746
- Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 922
- Configuring Aggregated Ethernet LACP (CLI Procedure) on page 926
- Understanding Aggregated Ethernet Interfaces and LACP on page 867
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

List of Sample Output **show lacp interfaces (Aggregated Ethernet) on page 1036**

Output Fields Table 141 on page 1033 lists the output fields for the **show lacp interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 141: show lacp interfaces Output Fields

Field Name	Field Description
Aggregated interface	Aggregated Ethernet interface value.

Table 141: show lacp interfaces Output Fields *(continued)*

Field Name	Field Description

LACP State	LACP state information for each aggregated Ethernet interface:
	• For a child interface configured with force-up, LACP state displays FUP along with the interface name.
	• **Role**—Role played by the interface. It can be one of the following:
	  ◦ **Actor**—Local device participating in LACP negotiation.
	  ◦ **Partner**—Remote device participating in LACP negotiation.
	• **Exp**—Expired state. **Yes** indicates the actor or partner is in an expired state. **No** indicates the actor or partner is not in an expired state.
	• **Def**—Default. **Yes** indicates that the actor's receive machine is using the default operational partner information, administratively configured for the partner. **No** indicates the operational partner information in use has been received in an LACP PDU.
	• **Dist**—Distribution of outgoing frames. **No** indicates distribution of outgoing frames on the link is currently disabled and is not expected to be enabled. Otherwise, the value is **Yes**.
	• **Col**—Collection of incoming frames. **Yes** indicates collection of incoming frames on the link is currently enabled and is not expected to be disabled. Otherwise, the value is **No**.
	• **Syn**—Synchronization. If the value is **Yes**, the link is considered synchronized. It has been allocated to the correct link aggregation group, the group has been associated with a compatible aggregator, and the identity of the link aggregation group is consistent with the system ID and operational key information transmitted. If the value is **No**, the link is not synchronized. It is currently not in the right aggregation.
	• **Aggr**—Ability of aggregation port to aggregate (**Yes**) or to operate only as an individual link (**No**).
	• **Timeout**—LACP timeout preference. Periodic transmissions of LACP PDUs occur at either a slow or fast transmission rate, depending upon the expressed LACP timeout preference (**Long Timeout** or **Short Timeout**).
	• **Activity**—Actor or partner's port activity. **Passive** indicates the port's preference for not transmitting LAC PDUs unless its partner's control value is **Active**. **Active** indicates the port's preference to participate in the protocol regardless of the partner's control value.

Table 141: show lacp interfaces Output Fields *(continued)*

Field Name	Field Description

LACP Protocol	LACP protocol information for each aggregated interface:
	• Link state (active or standby) indicated in parentheses next to the interface when link protection is configured.
	• **Receive State**—One of the following values:
	  • **Current**—The state machine receives an LACP PDU and enters the **Current** state.
	  • **Defaulted**—If no LACP PDU is received before the timer for the **Current** state expires a second time, the state machine enters the **Defaulted** state.
	  • **Expired**—If no LACP PDU is received before the timer for the **Current** state expires once, the state machine enters the **Expired** state.
	  • **Initialize**—When the physical connectivity of a link changes or a Begin event occurs, the state machine enters the **Initialize** state.
	  • **LACP Disabled**—If the port is operating in half duplex, the operation of LACP is disabled on the port, forcing the state to **LACP Disabled**. This state is similar to the **Defaulted** state, except that the port is forced to operate as an individual port.
	  • **Port Disabled**—If the port becomes inoperable and a Begin event has not occurred, the state machine enters the **Port Disabled** state.
	• **Transmit State**—Transmit state of state machine. One of the following values:
	  • **Fast Periodic**—Periodic transmissions are enabled at a fast transmission rate.
	  • **No Periodic**—Periodic transmissions are disabled.
	  • **Periodic Timer**—Transitory state entered when the periodic timer expires.
	  • **Slow Periodic**—Periodic transmissions are enabled at a slow transmission rate.
	• **Mux State**—State of the multiplexer state machine for the aggregation port. The state is one of the following values:
	  • **Attached**—Multiplexer state machine initiates the process of attaching the port to the selected aggregator.
	  • **Collecting**—**Yes** indicates that the receive function of this link is enabled with respect to its participation in an aggregation. Received frames are passed to the aggregator for collection. **No** indicates the receive function of this link is not enabled.
	  • **Collecting Distributing**—Collecting and distributing states are merged together to form a combined state (coupled control). Because independent control is not possible, the coupled control state machine does not wait for the partner to signal that collection has started before enabling both collection and distribution.
	  • **Detached**—Process of detaching the port from the aggregator is in progress.
	  • **Distributing**—**Yes** indicates that the transmit function of this link is enabled with respect to its participation in an aggregation. Frames may be passed down from the aggregator's distribution function for transmission. **No** indicates the transmit function of this link is not enabled.
	  • **Waiting**—Multiplexer state machine is in a holding process, awaiting an outcome.

LACP Statistics	LACP statistics are returned when the **extensive** option is used and provides the following information: • **LACP Rx**—LACP received counter that increments for each normal hello. • **LACP Tx**—Number of LACP transmit packet errors logged. • **Unknown Rx**—Number of unrecognized packet errors logged. • **Illegal Rx**—Number of invalid packets received.

show lacp interfaces (Aggregated Ethernet)

```
user@host> show lacp interfaces ae0 extensive
 Aggregated interface: ae0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout    Activity

      ge-1/0/1FUP    Actor    No   Yes    No   No   No   Yes     Fast      Active

      ge-1/0/1FUP  Partner    No   Yes    No   No   No   Yes     Fast     Passive

      ge-1/0/2       Actor    No   Yes    No   No   No   Yes     Fast      Active

      ge-1/0/2     Partner    No   Yes    No   No   No   Yes     Fast     Passive

    LACP protocol:        Receive State    Transmit State      Mux State
      ge-1/0/1FUP             CURRENT       Fast periodic      Collecting
distributing
      ge-1/0/2                CURRENT       Fast periodic      Collecting
distributing
      ge-1/0/1 (active)       CURRENT       Fast periodic      Collecting
distributing
      ge-1/0/2 (standby)      CURRENT       Fast periodic      WAITING
    LACP Statistics:       LACP Rx     LACP Tx   Unknown Rx   Illegal Rx
      ge-1/0/1                   0           0            0   0
      ge-1/0/2                   0           0            0   0
```

test interface restart-auto-negotiation

Syntax test interface restart-auto-negotiation *interface-name*

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Restarts auto-negotiation on a Fast Ethernet or Gigabit Ethernet interface.

Options *interface-name*—Interface name: **fe-*fpc/pic/port*** or **ge-*fpc/pic/port***.

Required Privilege Level view

List of Sample Output **test interface restart-auto-negotiation on page 1037**

Output Fields Use the **show interfaces extensive** command to see the state for auto-negotiation.

test interface restart-auto-negotiation

```
user@host> test interface restart-auto-negotiation fe-1/0/0
```

PART 13

Layer 2 Bridging and VLANs

CHAPTER 57

Bridging and VLANs—Overview

Understanding Bridging and VLANs on J-EX Series Switches

Network switches use Layer 2 bridging protocols to discover the topology of their LAN and to forward traffic toward destinations on the LAN.

This topic explains the following concepts regarding bridging and VLANs on J-EX Series Switches:

Ethernet LANs, Transparent Bridging, and VLANs

Ethernet is a data link layer technology, as defined by Layer 2 of the Open Systems Interconnection (OSI) model of communications protocols. Ethernet was first

standardized by the IEEE in 1982, in IEEE 802.3. Ethernet is used to create LANs. The network devices, called *nodes*, on the LAN transmit data in bundles that are generally called frames or packets.

Each node on a LAN has a unique identifier so that it can be unambiguously located on the network. Ethernet uses the Layer 2 media access control (MAC) address for this purpose. MAC addresses are hardware addresses that are programmed ("burned") into the Ethernet processor in the node.

A characteristic of Ethernet is that nodes on a LAN can transmit data frames at any time. However, the physical connecting cable between the nodes—either coaxial, copper-based (Category 5), or optical cable—can carry only a single stream of data at a time. One result of this design is that when two nodes transmit at the same time, their frames can collide on the cable and generate an error. Ethernet uses a protocol called carrier-sense multiple access with collision detection (CSMA/CD) to detect frame collisions. If a node receives a collision error message, it stops transmitting immediately and waits for a period of time before trying to send the frame again. If the node continues to detect collisions, it progressively increases the time between retransmissions in an attempt to find a time when no other data is being transmitted on the LAN. The node uses a backoff algorithm to calculate the increasing retransmission time intervals.

Ethernet LANs were originally implemented for small, simple networks that carried primarily text. Over time, LANs have become larger and more complex; the type of data they carry has grown to include voice, graphics, and video; and the increased speed of Ethernet interfaces on LANs has resulted in exponential increases in traffic on the network.

The IEEE 802.1D-2004 standard addresses some of the problems caused by the increase in LAN and complexity. This standard defines *transparent bridging* (generally called simply bridging). Bridging divides a single physical LAN (a single *broadcast domain*) into two or more virtual LANs, or VLANs. Each *VLAN* is a collection of network nodes that are grouped together to form separate broadcast domains. On an Ethernet network that is a single LAN, all traffic is forwarded to all nodes on the LAN. On VLANs, frames whose origin and destination are in the same VLAN are forwarded only within the local VLAN. Frames that are not destined for the local VLAN are the only ones forwarded to other broadcast domains. VLANs thus limit the amount of traffic flowing across the entire LAN, reducing the possible number of collisions and packet retransmissions within a VLAN and on the LAN as a whole.

On an Ethernet LAN, all network nodes must be physically connected to the same network. On VLANs, the physical location of the nodes is not important, so you can group network devices in any way that makes sense for your organization, such as by department or business function, types of network nodes, or even physical location. Each VLAN is identified by a single IP subnetwork and by standardized IEEE 802.1Q encapsulation (discussed below).

How Bridging Works

The transparent bridging protocol allows a switch to learn information about all the nodes on the LAN, including nodes on all the different VLANs. The switch uses this information to create address-lookup tables, called *Ethernet switching tables* that it consults when forwarding traffic to or toward a destination on the LAN.

Transparent bridging uses five mechanisms to create and maintain Ethernet switching tables on the switch:

- Learning
- Forwarding
- Flooding
- Filtering
- Aging

The first bridging mechanism is *learning*. When a switch is first connected to an Ethernet LAN or VLAN, it has no information about other nodes on the network. The switch goes through a learning process to obtain the MAC addresses of all the nodes on the network. It stores these in the Ethernet switching table. To learn MAC addresses, the switch reads all packets that it detects on the LAN or on the local VLAN, looking for MAC addresses of sending nodes. It places these addresses into its Ethernet switching table, along with two other pieces of information—the interface (or port) on which the traffic was received and the time when the address was learned.

The second bridging mechanism is *forwarding*. Switches forward traffic, passing it from an incoming interface to an outgoing interface that leads to or toward the destination. To forward frames, the switch consults the Ethernet switching table to see whether the table contains the MAC address corresponding to the frames' destination. If the Ethernet switching table contains an entry for the desired destination address, the switch sends the traffic out the interface associated with the MAC address. The switch also consults the Ethernet switching table in the same way when transmitting frames that originate on devices connected directly to the switch. If the Ethernet switching table does not contain an entry for the desired destination address, the switch uses flooding, which is the third bridging mechanism.

Flooding is how the switch learns about destinations not in its Ethernet switching table. If this table has no entry for a particular destination MAC address, the switch floods the traffic out all interfaces except the interface on which it was received. (If traffic originates on the switch, the switch floods it out all interfaces.) When the destination node receives the flooded traffic, it sends an acknowledgment packet back to the switch, allowing it to learn the MAC address of the node and to add the address to its Ethernet switching table.

Filtering, the fourth bridging mechanism, is how broadcast traffic is limited to the local VLAN whenever possible. As the number of entries in the Ethernet switching table grows, the switch pieces together an increasingly complete picture of the VLAN and the larger LAN—of which nodes are in the local VLAN and which are on other network segments. The switch uses this information to filter traffic. Specifically, for traffic whose source and destination MAC addresses are in the local VLAN, filtering prevents the switch from forwarding this traffic to other network segments.

Finally, the switch uses *aging*, the fifth bridging mechanism, to keep the entries in the Ethernet switching table current. For each MAC address in the Ethernet switching table, the switch records a timestamp of when the information about the network node was learned. Each time the switch detects traffic from a MAC address, it updates the

timestamp. A timer on the switch periodically checks the timestamp, and if it is older than a user-configured value, the switch removes the node's MAC address from the Ethernet switching table. This aging process ensures that the switch tracks only active nodes on the network and that it is able to flush out network nodes that are no longer available.

Types of Switch Ports

The ports, or interfaces, on a switch operate in either access mode or trunk mode.

An interface in access mode connects to a network device, such as a desktop computer, an IP telephone, a printer, a file server, or a security camera. The interface itself belongs to a single VLAN. The frames transmitted over an access interface are normal Ethernet frames. By default, when you boot a switch and use the factory-default configuration, or when you boot the switch and do not explicitly configure a port mode, all interfaces on the switch are in access mode.

Trunk interfaces handle traffic for multiple VLANs, multiplexing the traffic for all those VLANs over the same physical connection. Trunk interfaces are generally used to interconnect switches to one another.

IEEE 802.1Q Encapsulation and Tags

To identify which VLAN traffic belongs to, all frames on an Ethernet VLAN are identified by a tag, as defined in the IEEE 802.1Q standard. These frames are *tagged* and are encapsulated with 802.1Q tags.

For a simple network that has only a single VLAN, all traffic has the same 802.1Q tag.

When an Ethernet LAN is divided into VLANs, each VLAN is identified by a unique 802.1Q tag. The tag is applied to all frames so that the network nodes receiving the frames know which VLAN the frames belong to. Trunk ports, which multiplex traffic among a number of VLANs, use the tag to determine to origin of frames and where to forward them.

VLANs 0 and 4095 are reserved by the Junos OS, so you cannot use them in your network.

Assignment of Traffic to VLANs

You assign traffic to a particular VLAN in one of the following ways:

- By interface (port) on the switch. You specify that all traffic received on a particular interface on the switch is assigned to a specific VLAN. If you use the default factory switch settings, all traffic received on an access interface is untagged. This traffic is part of a default VLAN, but it is not tagged with an 802.1Q tag. When configuring the switch, you specify which VLAN to assign the traffic to. You configure the VLAN either by using a VLAN number (called a VLAN ID) or by using a name, which the switch translates into a numeric VLAN ID.
- By MAC address. You can specify that all traffic received from a specific MAC address be forwarded to a specific egress interface (next hop) on the switch. This method is administratively cumbersome to configure manually, but it can be useful when you are using automated databases to manage the switches on your network.

NOTE: If a J-EX4200 Ethernet Switch is interconnected with other switches in a Virtual Chassis configuration, each individual switch that is included as a member of the configuration is identified with a member ID. The member ID functions as an FPC slot number. When you are configuring interfaces for a Virtual Chassis configuration, you specify the appropriate member ID (0 through 9) as the *slot* element of the interface name.

The default factory settings for a Virtual Chassis configuration include FPC 0 as a member of the default VLAN because FPC 0 is configured as part of the ethernet-switching family. In order to include FPC 1 through FPC 9 in the default VLAN, add the ethernet-switching family to the configurations for those interfaces.

Ethernet Switching Tables

As J-EX Series switches learn the MAC addresses of the devices on local VLANs, they store them in the bridge on the switch. With each MAC address, the Ethernet switching table stores and associates the name of the interface (or port) on which the switch learned that address. The switch uses the information in this table when forwarding packets toward their destination.

Layer 2 and Layer 3 Forwarding of VLAN Traffic

To pass traffic within a VLAN, the switch uses Layer 2 forwarding protocols, including IEEE 802.1Q, Spanning Tree Protocol (STP), and GARP VLAN Registration Protocol (GVRP).

To pass traffic between two VLANs, the switch uses standard Layer 3 routing protocols, such as static routing, OSPF, and RIP. On J-EX Series switches, the same interfaces that support Layer 2 bridging protocols also support Layer 3 routing protocols, providing multilayer switching.

GVRP and MVRP

The GARP VLAN Registration Protocol (GVRP) and Multiple VLAN Registration Protocol (MVRP) are used to manage dynamic VLAN registration in a LAN.

GVRP is an application protocol of the Generic Attribute Registration Protocol (GARP) and is defined in the IEEE 802.1Q standard. GVRP learns VLANs on a particular 802.1Q trunk interface and adds the corresponding trunk interface to the VLAN if the advertised VLAN is preconfigured on the switch.

MVRP is an application protocol of the Multiple Registration Protocol (MRP) and is defined in the IEEE 802.1ak standard. MRP and MVRP were designed by IEEE to perform the same functions as GARP and GVRP while overcoming some GARP and GVRP limitations, in particular limitations involving bandwidth usage and convergence times in large networks with large numbers of VLANs. MVRP was created by IEEE as a replacement application for GVRP.

The VLAN registration information sent by MVRP and GVRP includes the current VLANs membership—that is, which switches are members of which VLANs—and which switch interfaces are in which VLAN. GVRP and MVRP share all VLAN information configured on a local switch.

MVRP can also be used to dynamically create VLANs, which are VLANs created on one switch and propagated to other switches as part of the MVRP message exchange process.

As part of ensuring that VLAN membership information is current, GVRP and MVRP remove switches and interfaces from the VLAN information when those switches and interfaces become unavailable. Pruning VLAN information has these benefits:

- Limits the network VLAN configuration to active participants only, reducing network overhead.
- Targets the scope of broadcast, unknown unicast, and multicast (BUM) traffic to interested devices only.

Routed VLAN Interface

In a traditional network, broadcast domains consist of either physical interfaces connected to a single switch or logical interfaces connected to one or more switches through VLAN configurations. Switches send traffic to hosts that are part of the same broadcast domain, but routers are needed to route traffic from one broadcast domain to another and to perform other Layer 3 functions such as traffic engineering. J-EX Series switches use a Layer 3 routed VLAN interface (RVI) named **vlan** to perform these routing functions, using it to route data to other Layer 3 interfaces. The RVI functions as a logical router, eliminating the need for having both a switch and a router.

The RVI (the **vlan** interface) must be configured as part of a broadcast domain or virtual private LAN service (VPLS) routing instance for Layer 3 traffic to be routed out of it. The RVI supports IPv4, IPv6, MPLS, and IS-IS traffic. At least one Layer 2 logical interface must be operational for the RVI to be operational. You must configure a broadcast domain or VPLS routing instance for the RVI just as you would configure a VLAN on the switch. Multicast data, broadcast data, or unicast data is switched between ports within the same RVI broadcast domain or VPLS routing instance. The RVI routes data that is destined for the switch's media access control (MAC) address.

Jumbo frames of up to 9216 bytes are supported on an RVI. To route jumbo data packets on the RVI, you must configure the jumbo MTU size on the member physical interfaces of the RVI and not on the RVI itself (the **vlan** interface). However, for jumbo control packets—for example, to ping the RVI with a packet size of 6000 bytes or more—you must explicitly configure the jumbo MTU size on the interface named **vlan** (the RVI).

CAUTION: Setting or deleting the jumbo MTU size on the RVI (the vlan interface) while the switch is transmitting packets might result in dropped packets.

See "Configuring Routed VLAN Interfaces (CLI Procedure)" on page 1137.

To learn more about configuring routing protocols and policies, see the *Junos OS Routing Protocols Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.

Related Documentation

- Understanding Layer 2 Protocol Tunneling on J-EX Series Switches on page 1056
- Understanding Multiple VLAN Registration Protocol (MVRP) on J-EX Series Switches on page 1054
- Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch on page 1063
- Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches on page 1070
- Example: Configure Automatic VLAN Administration Using GVRP on page 1087
- Example: Connecting an Access Switch to a Distribution Switch on page 1078

Understanding Private VLANs on J-EX Series Switches

The private VLAN (PVLAN) feature on J-EX Series Switches allows an administrator to split a broadcast domain into multiple isolated broadcast subdomains, essentially putting a VLAN inside a VLAN. Just like regular VLANs, PVLANs are isolated on Layer 2 and require that a Layer 3 device be used to route traffic among them. Private VLANs are useful for restricting the flow of broadcast and unknown unicast traffic and for limiting the communication between known hosts.

NOTE: Configuring a voice over IP (VoIP) VLAN on PVLAN interfaces is not supported.

In a private VLAN, one VLAN is designated the primary VLAN, and other VLANs are nested inside that VLAN as secondary VLANs.

- Primary—A VLAN used to forward frames downstream to isolated and community VLANs.
- Isolated—A secondary VLAN that receives packets only from the primary VLAN and forwards frames upstream to the primary VLAN.
- Community—A secondary VLAN that transports frames among community interfaces within the same community and forwards frames upstream to the primary VLAN.

Private VLANs provide IP address conservation and efficient allocation of those IP addresses. In a typical network, VLANs usually correspond to a single IP subnet. In private VLANs, the hosts in all the secondary VLANs still belong to the same IP subnet as the subnet allocated to the primary VLAN. Hosts within the secondary VLAN are numbered out of IP subnets associated with the primary VLAN, and their IP subnet masking information reflects that of the primary VLAN subnet. Any primary routed VLAN interfaces (RVIs) perform functions similar to proxy ARP to enable communication between hosts that are members of a different secondary VLAN.

NOTE: If you enable no-mac-learning on a primary VLAN, all isolated VLANs in that private VLAN inherit that setting. If you want to disable MAC address learning on any community VLANs, you must configure no-mac-learning on each of those VLANs.

Related Documentation

- Understanding Bridging and VLANs on J-EX Series Switches on page 1041
- Example: Configuring a Private VLAN on a J-EX Series Switch on page 1107
- Creating a Private VLAN (CLI Procedure) on page 1143

Understanding Virtual Routing Instances on J-EX Series Switches

Virtual routing instances allow administrators to divide a J-EX Series Switch into multiple independent virtual routers, each with its own routing table. Splitting a device into many virtual routing instances isolates traffic traveling across the network without requiring multiple devices to segment the network.

You can use virtual routing instances to isolate customer traffic on your network and to bind customer-specific instances to customer-owned interfaces.

Virtual routing and forwarding (VRF) is often used in conjunction with Layer 3 subinterfaces, allowing traffic on a single physical interface to be differentiated and associated with multiple virtual routers. Each logical Layer 3 subinterface can belong to only one routing instance.

J-EX Series switches support IPv4 and IPv6 unicast and multicast VRF traffic.

J-EX4200 Series Ethernet Switches support up to 252 IPv4 virtual routing instances and up to 64 IPv6 virtual routing instances. J-EX8200 Series Ethernet Switches support up to 252 IPv4 and IPv6 virtual routing instances.

Related Documentation

- Understanding Layer 3 Subinterfaces on page 871
- Example: Using Virtual Routing Instances to Route Among VLANs on J-EX Series Switches on page 1112
- Configuring Virtual Routing Instances (CLI Procedure) on page 1142

Understanding Redundant Trunk Links on J-EX Series Switches

In a typical enterprise network comprised of distribution and access layers, a redundant trunk link provides a simple solution for network recovery when a trunk port goes down. Traffic is routed to another trunk port, keeping network convergence time to a minimum. You can configure a maximum of 16 redundant trunk groups on a standalone switch or on a Virtual Chassis.

To configure a redundant trunk link, create a redundant trunk group. The redundant trunk group is configured on the access switch, and contains two links: a primary or active link, and a secondary link. If the active link fails, the secondary link automatically starts forwarding data traffic without waiting for normal STP convergence.

Data traffic is forwarded only on the active link. Data Traffic on the secondary link is dropped and shown as dropped packets when you issue the operational mode command **show interfaces xe-** ***interface-name*** **extensive**.

While data traffic is blocked on the secondary link, Layer 2 control traffic is still permitted. For example, an LLDP session can be run between two J-EX Series Switches on the secondary link.

STP is enabled by default on J-EX Series switches to create a loop-free topology. When trunk links are placed in a redundant group, they cannot be part of an STP topology. The Junos OS for J-EX Series switches does not allow an interface to be in a redundant trunk group and in an STP topology at the same time. However, STP can continue operating in other parts of the network. For example, STP may continue operating between the distribution switches and linking them to the enterprise core.

Figure 26 on page 1050 shows three switches in a basic topology for redundant trunk links. Switch 1 and Switch 2 make up the distribution layer, and Switch 3 makes up the access layer. Switch 3 is connected to the distribution layer through trunk ports **ge-0/0/9.0** (Link 1) and **ge-0/0/10.0** (Link 2). Link 1 and Link 2 are in a redundant trunk group called **group1**. Link 1 is designated as the primary link. Traffic flows between Switch 3 in the access layer and Switch 1 in the distribution layer through Link 1. While Link 1 is active, Link 2 blocks traffic.

Figure 26: Redundant Trunk Group, Link 1 Active

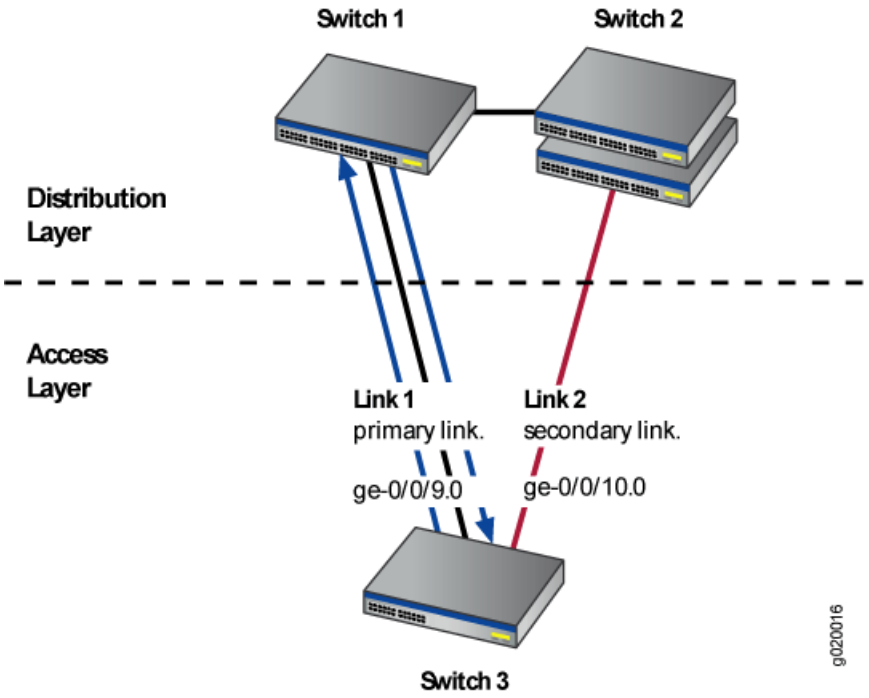

Figure 27 on page 1050 illustrates how the redundant trunk link topology works when the primary link goes down.

Figure 27: Redundant Trunk Group, Link 2 Active

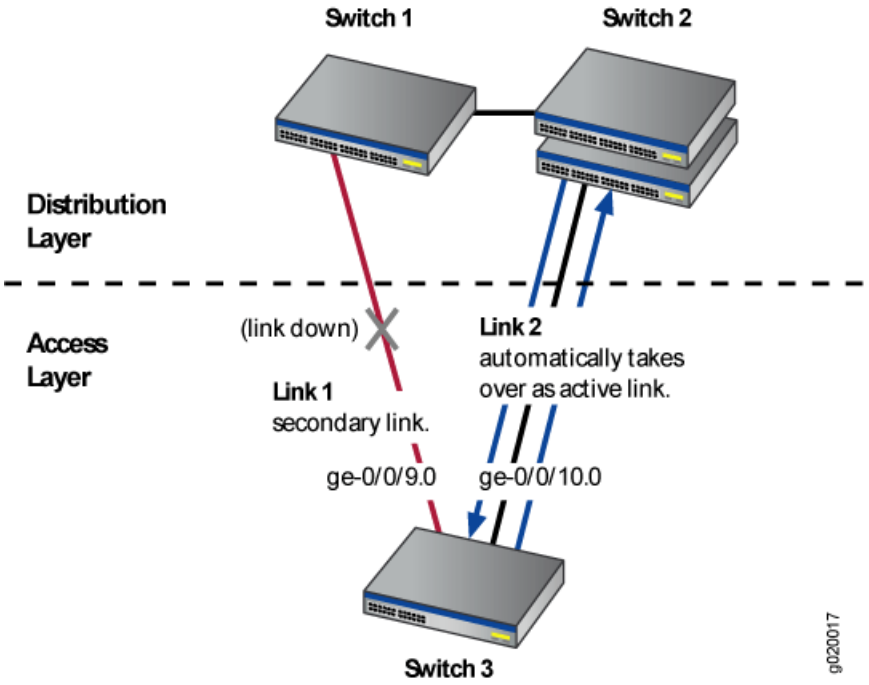

Link 1 is down between Switch 3 and Switch 1. Link 2 takes over as the active link. Traffic between the access layer and the distribution layer is automatically switched to Link 2 between Switch 1 and Switch 2.

Related Documentation

- Example: Configuring Redundant Trunk Links for Faster Recovery on page 1101
- **redundant-trunk-group on page 1218**

Understanding Q-in-Q Tunneling on J-EX Series Switches

Q-in-Q tunneling allows service providers on Ethernet access networks to extend a Layer 2 Ethernet connection between two customer sites. Using Q-in-Q tunneling, providers can also segregate or bundle customer traffic into fewer VLANs or different VLANs by adding another layer of 802.1Q tags. Q-in-Q tunneling is useful when customers have overlapping VLAN IDs, because the customer's 802.1Q (dot1Q) VLAN tags are prepended by the service VLAN (S-VLAN) tag. The Junos OS implementation of Q-in-Q tunneling supports the IEEE 802.1ad standard.

This topic describes:

- How Q-in-Q Tunneling Works on page 1051
- Disabling MAC Address Learning on page 1052
- Mapping C-VLANs to S-VLANs on page 1052
- Routed VLAN Interfaces on Q-in-Q VLANs on page 1053
- Limitations for Q-in-Q Tunneling on page 1054

How Q-in-Q Tunneling Works

In Q-in-Q tunneling, as a packet travels from a customer VLAN (C-VLAN) to a service provider's VLAN, a customer-specific 802.1Q tag is added to the packet. This additional tag is used to segregate traffic into service-provider-defined service VLANs (S-VLANs). The original customer 802.1Q tag of the packet remains and is transmitted transparently, passing through the service provider's network. As the packet leaves the S-VLAN in the downstream direction, the extra 802.1Q tag is removed.

When Q-in-Q tunneling is enabled on J-EX Series Switches, trunk interfaces are assumed to be part of the service provider network and access interfaces are assumed to be customer facing. An access interface can receive both tagged and untagged frames in this case.

An interface can be a member of multiple S-VLANs. You can map one C-VLAN to one S-VLAN (1:1) or multiple C-VLANs to one S-VLAN (N:1). Packets are double-tagged for an additional layer of segregating or bundling of C-VLANs. C-VLAN and S-VLAN tags are unique; so you can have both a C-VLAN 101 and an S-VLAN 101, for example. You can limit the set of accepted customer tags to a range of tags or to discrete values. Class-of-service (CoS) values of C-VLANs are unchanged in the downstream direction. You may, optionally, copy ingress priority and CoS settings to the S-VLAN. Using private VLANs, you can isolate users to prevent the forwarding of traffic between user interfaces even if the interfaces are on the same VLAN.

You can use the **native** option to specify an S-VLAN for untagged and priority tagged packets when using many-to-one bundling and mapping a specific interface approaches to map C-VLANs to S-VLANs. Otherwise the packets are discarded. The **native** option is not available for all-in-one bundling because there is no need to specify untagged and priority tagged packets when all packets are mapped to the C-VLAN. See the Mapping C-VLANs to S-VLANs section of this document for information on the methods of mapping C-VLANs to S-VLANs.

Firewall filters allow you to map an interface to a VLAN based on a policy. Using firewall filters to map an interface to a VLAN is useful when you want a subset of traffic from a port to be mapped to a selected VLAN instead of the designated VLAN. To configure a firewall filter to map an interface to a VLAN, the **vlan** option has to be configured as part of the firewall filter and the **mapping policy** option must be specified in the interface configuration for each logical interface using the filter.

Disabling MAC Address Learning

In a Q-in-Q deployment, customer packets from downstream interfaces are transported without any changes to source and destination MAC addresses. You can disable MAC address learning at both the interface level and the VLAN level. Disabling MAC address learning on an interface disables learning for all the VLANs of which that interface is a member. When you disable MAC address learning on a VLAN, MAC addresses that have already been learned are flushed.

If you disable MAC address learning on an interface or a VLAN, you cannot include MAC move limiting or 802.1X authentication in that same VLAN configuration.

When a routed VLAN interface (RVI) is associated with either an interface or a VLAN on which MAC address learning is disabled, the Layer 3 routes resolved on that VLAN or that interface are not resolved with the Layer 2 component. This results in routed packets flooding all the interfaces associated with the VLAN.

Mapping C-VLANs to S-VLANs

There are three ways to map C-VLANs to an S-VLAN:

- All-in-one bundling—Use the **dot1q-tunneling** option to map without specifying customer VLANs. All packets from all access interfaces are mapped to the S-VLAN.
- Many-to-one bundling—Use the **customer-vlans** option to specify which C-VLANs are mapped to the S-VLAN.
- Mapping a specific interface—Use the **mapping** option to indicate a specific S-VLAN for a given C-VLAN. The specified C-VLAN applies to only one VLAN and not all access interfaces as in the cases of all-in-one and many-to-one bundling.

If you configure multiple methods, the switch gives priority to mapping a specific interface, then to many-to-one bundling, and last to all-in-one bundling. However, you cannot have overlapping rules for the same C-VLAN under a given approach.

- All-in-One Bundling on page 1053
- Many-to-One Bundling on page 1053
- Mapping a Specific Interface on page 1053

All-in-One Bundling

All-in-one bundling maps all packets from all access interfaces to the S-VLAN. All-in-one bundling is configured using the **dot1q-tunneling** option without specifying customer VLANs.

When all-in-one bundling is used, all packets leaving the C-VLAN, including untagged and priority tagged packets, enter the S-VLAN.

Many-to-One Bundling

Many-to-one bundling is used to specify which C-VLANs are mapped to an S-VLAN. Many-to-one bundling is configured using the **customer-vlans** option.

Many-to-one bundling is used when you want a subset of the C-VLANs on the access switch to be part of the S-VLAN. When using many-to-one bundling, untagged and priority tagged packets can be mapped to the S-VLAN when the **native** option is specified along with the **customer-vlans** option.

Mapping a Specific Interface

Use the mapping a specific interface approach when you want to assign an S-VLAN to a specific C-VLAN on an interface. The mapping a specific interface configuration only applies to the configured interface, not to all access interfaces as in the cases of the all-in-one bundling and many-to-one bundling approaches. The mapping a specific interface approach is configured using the **mapping** option to indicate a specific S-VLAN for a given C-VLAN.

The mapping a specific interface approach has two suboptions for treatment of traffic: swap and push. When traffic that is mapped to a specific interface is pushed, the packet retains its tag as it moves between the S-VLAN and C-VLAN and an additional VLAN tag is added to the packet. When traffic that is mapped to a specific interface is swapped, the incoming tag is replaced with a new VLAN tag. Using the **swap** option is also referred to as VLAN ID translation.

It might be useful to have S-VLANs that provide service to multiple customers. Each customer will typically have its own S-VLAN plus access to one or more S-VLANs that are used by multiple customers. A specific tag on the customer side is mapped to an S-VLAN. Typically, this functionality is used to keep data from different customers separate or to provide individualized treatment of the packets on a certain interface.

Routed VLAN Interfaces on Q-in-Q VLANs

Routed VLAN interfaces (RVIs) are supported on Q-in-Q VLANs.

Packets arriving on an RVI that is using Q-in-Q VLANs will get routed regardless of whether the packet is single or double tagged. The outgoing routed packets contain an S-VLAN tag only when exiting a trunk interface; the packets exit the interface untagged when exiting an access interface.

Limitations for Q-in-Q Tunneling

Q-in-Q tunneling does not support most access port security features. There is no per-VLAN (customer) policing or per-VLAN (outgoing) shaping and limiting with Q-in-Q tunneling unless you configure these security features using firewall filters.

Related Documentation

- Understanding Bridging and VLANs on J-EX Series Switches on page 1041
- Example: Setting Up Q-in-Q Tunneling on J-EX Series Switches on page 1105
- Configuring Q-in-Q Tunneling (CLI Procedure) on page 1144

Understanding Multiple VLAN Registration Protocol (MVRP) on J-EX Series Switches

You can configure Multiple VLAN Registration Protocol (MVRP) on J-EX Series Switches. The primary purpose of MVRP is to manage dynamic VLAN registration in a LAN. In managing dynamic VLAN registration, MVRP also prunes VLAN information. MVRP can also be used to dynamically create VLANs in switching networks.

MVRP is an application protocol of the Multiple Registration Protocol (MRP) and is defined in the IEEE 802.1ak standard. MRP and MVRP were designed by IEEE to perform the same functions as Generic Attribute Registration Protocol (GARP) and GARP VLAN Registration Protocol (GVRP) while overcoming some GARP and GVRP limitations, in particular limitations involving bandwidth usage and convergence time in large networks with large numbers of VLANs.

MVRP was created by IEEE as a replacement application for GVRP. MVRP and GVRP cannot be run concurrently to share VLAN information in a switching network.

This topic describes:

- How MVRP Works on J-EX Series Switches on page 1054
- Basics of MVRP on J-EX Series Switches on page 1055
- MVRP Registration Modes on page 1055
- MRP Timers on page 1055
- MRP VLAN Messages on page 1056

How MVRP Works on J-EX Series Switches

The VLAN registration information sent by MVRP protocol data units (PDUs) includes the current VLANs membership—that is, which switches are members of which VLANs—and which switch interfaces are in which VLAN. MVRP shares all information in the PDU with all switches participating in MVRP in the switching network.

MVRP stays synchronized using these PDUs. The MVRP PDUs are sent to other switches on the network only when an MVRP state change occurs. The switches in the network

participating in MVRP receive these PDUs during state changes and update their MVRP states accordingly. MVRP timers dictate when PDUs can be sent and when switches receiving MVRP PDUs can update their MVRP information.

VLAN information is distributed as part of the MVRP message exchange process and can be used to dynamically create VLANs, which are VLANs created on one switch and propagated to other switches as part of the MVRP message exchange process. Dynamic VLAN creation using MVRP is enabled by default but can be disabled.

As part of ensuring that VLAN membership information is current, MVRP removes switches and interfaces from the VLAN information when they become unavailable. Pruning VLAN information has these benefits:

- Limits the network VLAN configuration to active participants only, reducing network overhead.
- Targets the scope of broadcast, unknown unicast, and multicast (BUM) traffic to interested devices only.

Basics of MVRP on J-EX Series Switches

MVRP is disabled by default on all J-EX Series switches. You can configure MVRP on J-EX Series switch interfaces to participate in MVRP for the switching network. MVRP can only be enabled on trunk interfaces, and dynamic VLAN configuration through MVRP is enabled by default when MVRP is enabled.

MVRP Registration Modes

The MVRP registration mode defines whether an interface does or does not participate in MVRP.

The following MVRP registration modes are configurable:

- forbidden—The interface does not register and does not participate in MVRP.
- normal—The interface accepts MVRP messages and participates in MVRP. This is the default registration mode setting.

MRP Timers

MVRP registration and updates are controlled by timers that are part of the MRP protocol. These timers are set on a per-interface basis and define when MVRP PDUs can be sent and when MVRP information can be updated on a switch.

The following timers are used to control the operation of MVRP:

- Join timer—Controls the interval for the next MVRP PDU transmit opportunity.
- Leave timer—Controls the period of time that an interface on the switch waits in the Leave state before changing to the unregistered state.
- LeaveAll timer—Controls the frequency with which the interface generates LeaveAll messages.

BEST PRACTICE: Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.

MRP VLAN Messages

MVRP uses MRP messages to register and declare MVRP states for a switch and to inform the switching network that a switch is leaving MVRP. These messages are communicated as part of the PDU to communicate the state of a particular switch interface on the switching network to the other switches in the network.

The following messages are communicated for MVRP:

- Empty—VLAN information is not being declared and is not registered.
- In—VLAN information is not being declared but is registered.
- JoinEmpty—VLAN information is being declared but not registered.
- JoinIn—VLAN information is being declared and is registered.
- Leave—VLAN information that was previously registered is being withdrawn.
- LeaveAll—All registrations will be de-registered. Participants that want to participate in MVRP will need to re-register.
- New—VLAN information is new and possibly not previously registered.

Related Documentation

- Understanding Bridging and VLANs on J-EX Series Switches on page 1041
- Example: Configuring Automatic VLAN Administration Using MVRP on J-EX Series Switches on page 1115
- Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure) on page 1147

Understanding Layer 2 Protocol Tunneling on J-EX Series Switches

Layer 2 protocol tunneling (L2PT) allows you to send Layer 2 protocol data units (PDUs) across a service provider network and deliver them to J-EX Series Switches that are not part of the local broadcast domain. This feature is useful when you want to run Layer 2 protocols on a network that includes switches located at remote sites that are connected across a service provider network.

This topic includes:

- Layer 2 Protocols Supported by L2PT on J-EX Series Switches on page 1057
- How L2PT Works on page 1057
- L2PT Basics on J-EX Series Switches on page 1058

Layer 2 Protocols Supported by L2PT on J-EX Series Switches

L2PT on J-EX Series switches supports the following Layer 2 protocols:

- 802.1X authentication
- 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM)

NOTE: If you enable L2PT for untagged OAM LFM packets, do not configure LFM on the corresponding access interface.

- Cisco Discovery Protocol (CDP)
- Ethernet local management interface (E-LMI)
- GARP VLAN Registration Protocol (GVRP)
- Link Aggregation Control Protocol (LACP)

NOTE: If you enable L2PT for untagged LACP packets, do not configure LACP on the corresponding access interface.

- Link Layer Discovery Protocol (LLDP)
- Multiple MAC Registration Protocol (MMRP)
- Multiple VLAN Registration Protocol (MVRP)
- Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP)
- VLAN Spanning Tree Protocol (VSTP)
- VLAN Trunking Protocol (VTP)

NOTE: CDP and VTP cannot be configured on J-EX Series switches. L2PT does, however, tunnel CDP and VTP PDUs.

How L2PT Works

L2PT works by encapsulating Layer 2 PDUs, tunneling them across a service provider network, and decapsulating them for delivery to their destination switches. L2PT encapsulates Layer 2 PDUs by enabling the ingress provider edge (PE) device to rewrite the PDUs' destination media access control (MAC) addresses before forwarding them onto the service provider network. The devices in the service provider network treat these encapsulated PDUs as multicast Ethernet packets. Upon receipt of these PDUs, the egress PE devices decapsulate them by replacing the destination MAC addresses with the address of the Layer 2 protocol that is being tunneled before forwarding the PDUs to their destination switches.

L2PT Basics on J-EX Series Switches

L2PT is enabled on a per-VLAN basis. When you enable L2PT on a VLAN, all access interfaces are considered to be customer-facing interfaces, all trunk interfaces are considered to be service provider network-facing interfaces, and the specified Layer 2 protocol is disabled on the access interfaces. L2PT only acts on logical interfaces of the family **ethernet-switching**.

NOTE: Access interfaces in an L2PT-enabled VLAN should not receive L2PT-tunneled PDUs. If an access interface does receive L2PT-tunneled PDUs, it might mean that there is a loop in the network. As a result, the interface will be shut down.

L2PT is configured under the **[edit vlans *vlan-name* dot1q-tunneling]** hierarchy level, meaning Q-in-Q tunneling is (and must be) enabled. If L2PT is not enabled, Layer 2 PDUs are handled in the same way they were handled before L2PT was enabled.

NOTE: If the switch receives untagged or priority-tagged Layer 2 control PDUs to be tunnelled, then you must configure the switch to map untagged and priority-tagged packets to an L2PT-enabled VLAN. For more information on assigning untagged and priority-tagged packets to VLANs, see “Understanding Q-in-Q Tunneling on J-EX Series Switches” on page 1051 and “Configuring Q-in-Q Tunneling (CLI Procedure)” on page 1144.

Related Documentation

- Example: Configuring Layer 2 Protocol Tunneling on J-EX Series Switches on page 1126
- Example: Setting Up Q-in-Q Tunneling on J-EX Series Switches on page 1105

Understanding Proxy ARP on EX Series Switches

You can configure proxy Address Resolution Protocol (ARP) on your J-EX Series Ethernet switch to enable the switch to respond to ARP queries for network addresses by offering its own Ethernet media access control (MAC) address. With proxy ARP enabled, the switch captures and routes traffic to the intended destination.

Proxy ARP is useful in situations where hosts are on different physical networks and you do not want to use subnet masking. Because ARP broadcasts are not propagated between hosts on different physical networks, hosts will not receive a response to their ARP request if the destination is on a different subnet. Enabling the switch to act as an ARP proxy allows the hosts to transparently communicate with each other through the switch. Proxy ARP can help hosts on a subnet reach remote subnets without your having to configure routing or a default gateway.

What Is ARP?

Ethernet LANs use ARP to map Ethernet MAC addresses to IP addresses. Each device maintains a cache containing a mapping of MAC addresses to IP addresses. The switch maintains this mapping in a cache that it consults when forwarding packets to network devices. If the ARP cache does not contain an entry for the destination device, the host (the DHCP client) broadcasts an ARP request for that device's address and stores the response in the cache.

Proxy ARP Overview

When proxy ARP is enabled, if the switch receives an ARP request for which it has a route to the target (destination) IP address, the switch responds by sending a proxy ARP reply packet containing its own MAC address. The host that sent the ARP request then sends its packets to the switch, which forwards them to the intended host.

NOTE: For security reasons, the source address in an ARP request must be on the same subnet as the interface on which the ARP request is received.

You can configure proxy ARP for each interface. You can also configure proxy ARP for a VLAN by using a routed VLAN interface (RVI).

J-EX Series switches support two modes of proxy ARP, restricted and unrestricted. Both modes require that the switch have an active route to the destination address of the ARP request.

- Restricted—The switch responds to ARP requests in which the physical networks of the source and target are different and does not respond if the source and target IP addresses are on the same subnet. In this mode, hosts on the same subnet communicate without proxy ARP. We recommend that you use this mode on the switch.

- Unrestricted—The switch responds to all ARP requests for which it has a route to the destination. This is the default mode (because it is the default mode in the Junos operating system (Junos OS) configurations other than those on the switch). We recommend using restricted mode on the switch.

Best Practices for Proxy ARP on J-EX Series Switches

We recommend these best practices for configuring proxy ARP on the switches:

- Set proxy ARP to restricted mode.
- Use restricted mode when configuring proxy ARP on RVIs.
- If you set proxy ARP to unrestricted, disable gratuitous ARP requests on each interface enabled for proxy ARP.

Related Documentation

- Example: Configuring Proxy ARP on a J-EX Series Switch on page 2621
- Configuring Proxy ARP (CLI Procedure) on page 1153

Understanding MAC Notification on J-EX Series Switches

J-EX Series Switches track clients on a network by storing Media Access Control (MAC) addresses in the Ethernet switching table on the switch. When switches learn or unlearn a MAC address, SNMP notifications can be sent to the network management system at regular intervals to record the addition or removal of the MAC address. This process is known as MAC notification.

The MAC Notification MIB controls MAC notification for the network management system. For general information on the MAC Notification MIB, see the *Junos OS Network Management Configuration Guide* at http://www.juniper.net/techpubs/software/junos/index.html.

The MAC notification interval defines how often these SNMP notifications are sent to the network management system. The MAC notification interval works by tracking all of the MAC address additions or removals on the switch over a period of time and then sending all of the tracked MAC address additions or removals to the network management server at the end of the interval. For instance, if the MAC notification interval is set to 10, all of the MAC address addition and removal SNMP notifications are sent to the network management system every 10 seconds.

Enabling MAC notification allows users to monitor the addition and removal of MAC addresses from the Ethernet switching table remotely using a network management system. The advantage of setting a high MAC notification interval is that the amount of network traffic is reduced because updates are sent less frequently. The advantage of setting a low MAC notification interval is that the network management system is better synchronized with the switch.

MAC notification is disabled by default. When MAC notification is enabled, the default MAC notification interval is 30 seconds.

Related Documentation

- Configuring MAC Notification (CLI Procedure) on page 1151
- Configuring SNMP (J-Web Procedure) on page 3309

CHAPTER 58

Examples: Bridging and VLAN Configuration

Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch

J-EX Series switches use bridging and virtual LANs (VLANs) to connect network devices in a LAN—desktop computers, IP telephones, printers, file servers, wireless access points, and others—and to segment the LAN into smaller bridging domains. The switch's default configuration provides a quick setup of bridging and a single VLAN.

This example describes how to configure basic bridging and VLANs for a J-EX Series switch:

Requirements

This example uses the following software and hardware components:

- One J-EX4200 Virtual Chassis switch

Before you set up bridging and a VLAN, be sure you have:

- Installed your J-EX Series switch. See Installing and Connecting a J-EX4200 Switch.
- Performed the initial switch configuration. See "Connecting and Configuring a J-EX Series Switch (J-Web Procedure)" on page 163.

Overview and Topology

J-EX Series switches connect network devices in an office LAN or a data center LAN to provide sharing of common resources such as printers and file servers and to enable wireless devices to connect to the LAN through wireless access points. Without bridging and VLANs, all devices on the Ethernet LAN are in a single broadcast domain, and all the devices detect all the packets on the LAN. Bridging creates separate broadcast domains on the LAN, creating VLANs, which are independent logical networks that group together related devices into separate network segments. The grouping of devices on a VLAN is independent of where the devices are physically located in the LAN.

To use a J-EX Series switch to connect network devices on a LAN, you must, at a minimum, configure bridging and VLANs. If you simply power on the switch and perform the initial switch configuration using the factory-default settings, bridging is enabled on all the switch's interfaces, all interfaces are in access mode, and all interfaces belong to a VLAN called **default**, which is automatically configured. When you plug access devices—such as desktop computers, Avaya IP telephones, file servers, printers, and wireless access points—into the switch, they are joined immediately into the **default** VLAN and the LAN is up and running.

The topology used in this example consists of one J-EX4200-24T switch, which has a total of 24 ports. Eight of the ports support Power over Ethernet (PoE), which means they provide both network connectivity and electric power for the device connecting to the port. To these ports, you can plug in devices requiring PoE, such as Avaya VoIP telephones, wireless access points, and some IP cameras. (Avaya phones have a built-in hub that allows you to connect a desktop PC to the phone, so the desktop and phone in a single office require only one port on the switch.) The remaining 16 ports provide only network connectivity. You use them to connect devices that have their own power sources, such as desktop and laptop computers, printers, and servers. Table 1 details the topology used in this configuration example.

Table 142: Components of the Basic Bridging Configuration Topology

Property	Settings
Switch hardware	J-EX4200-24T switch, with 24 Gigabit Ethernet ports: 8 PoE ports (**ge-0/0/0** through **ge-0/0/7**) and 16 non-PoE ports (**ge-0/0/8** through **ge-0/0/23**)
VLAN name	**default**
Connection to wireless access point (requires PoE)	**ge-0/0/0**

Table 142: Components of the Basic Bridging Configuration Topology *(continued)*

Property	Settings
Connections to Avaya IP telephone—with integrated hub, to connect phone and desktop PC to a single port (requires PoE)	**ge-0/0/1** through **ge-0/0/7**
Direct connections to desktop PCs (no PoE required)	**ge-0/0/8** through **ge-0/0/12**
Connections to file servers (no PoE required)	**ge-0/0/17** and **ge-0/0/18**
Connections to integrated printer/fax/copier machines (no PoE required)	**ge-0/0/19** through **ge-0/0/20**
Unused ports (for future expansion)	**ge-0/0/13** through **ge-0/0/16**, and **ge-0/0/21** through **ge-0/0/23**

Configuration

CLI Quick Configuration

By default, after you perform the initial configuration on the J-EX4200 switch, switching is enabled on all interfaces, a VLAN named **default** is created, and all interfaces are placed into this VLAN. You do not need to perform any other configuration on the switch to set up bridging and VLANs. To use the switch, simply plug the Avaya IP phones into the PoE-enabled ports **ge-0/0/1** through **ge-0/0/7**, and plug in the PCs, file servers, and printers to the non-PoE ports, **ge-0/0/8** through **ge-0/0/12** and **ge-0/0/17** through **ge-0/0/20**.

Step-by-Step Procedure

To configure bridging and VLANs:

1. Make sure the switch is powered on.
2. Connect the wireless access point to switch port **ge-0/0/0**.
3. Connect the seven Avaya phones to switch ports **ge-0/0/1** through **ge-0/0/7**.
4. Connect the five PCs to ports **ge-0/0/8** through **ge-0/0/12**.
5. Connect the two file servers to ports **ge-0/0/17** and **ge-0/0/18**.
6. Connect the two printers to ports **ge-0/0/19** and **ge-0/0/20**.

Results

Check the results of the configuration:

```
[edit]
user@switch> show configuration
## Last commit: 2008-03-06 00:11:22 UTC by triumph
version 9.0;
system {
   root-authentication {
      encrypted-password "$1$urmA7AFM$x5SaGEUOdSI3u1K/iITGh1"; ## SECRET-DATA
   }
   syslog {
      user * {
         any emergency;
```

```
        }
        file messages {
            any notice;
            authorization info;
        }
        file interactive-commands {
            interactive-commands any;
        }
    }
    commit {
        factory-settings {
            reset-chassis-lcd-menu;
            reset-virtual-chassis-configuration;
        }
    }
}
interfaces {
    ge-0/0/0 {
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/0/1 {
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/0/2 {
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/0/3 {
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/0/4 {
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/0/5 {
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/0/6 {
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/0/7 {
        unit 0 {
            family ethernet-switching;
        }
```

```
}
ge-0/0/8 {
    unit 0 {
        family ethernet-switching;
    }
}
ge-0/0/9 {
    unit 0 {
        family ethernet-switching;
    }
}
ge-0/0/10 {
    unit 0 {
        family ethernet-switching;
    }
}
ge-0/0/11 {
    unit 0 {
        family ethernet-switching;
    }
}
ge-0/0/12 {
    unit 0 {
        family ethernet-switching;
    }
}
ge-0/0/13 {
    unit 0 {
        family ethernet-switching;
    }
}
ge-0/0/14 {
    unit 0 {
        family ethernet-switching;
    }
}
ge-0/0/15 {
    unit 0 {
        family ethernet-switching;
    }
}
ge-0/0/16 {
    unit 0 {
        family ethernet-switching;
    }
}
ge-0/0/17 {
    unit 0 {
        family ethernet-switching;
    }
}
ge-0/0/18 {
    unit 0 {
        family ethernet-switching;
    }
}
```

```
    ge-0/0/19 {
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/0/20 {
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/0/21 {
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/0/22 {
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/0/23 {
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/1/0 {
        unit 0 {
            family ethernet-switching;
        }
    }
    xe-0/1/0 {
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/1/1 {
        unit 0 {
            family ethernet-switching;
        }
    }
    xe-0/1/1 {
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/1/2 {
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/1/3 {
        unit 0 {
            family ethernet-switching;
        }
    }
}
```

```
protocols {
  lldp {
    interface all;
  }
  rstp;
}
  poe {
    interface all;
  }
```

Verification

To verify that switching is operational and that a VLAN has been created, perform these tasks:

- Verifying That the VLAN Has Been Created on page 1069
- Verifying That Interfaces Are Associated with the Proper VLANs on page 1069

Verifying That the VLAN Has Been Created

Purpose Verify that the VLAN named **default** has been created on the switch.

Action List all VLANs configured on the switch:

```
user@switch> show vlans

Name            Tag     Interfaces
default
                        ge-0/0/0.0*, ge-0/0/1.0, ge-0/0/2.0, ge-0/0/3.0,
                        ge-0/0/4.0,  ge-0/0/5.0, ge-0/0/6.0, ge-0/0/7.0,
                        ge-0/0/8.0*, ge-0/0/9.0, ge-0/0/10.0, ge-0/0/11.0*,
                        ge-0/0/12.0, ge-0/0/13.0, ge-0/0/14.0, ge-0/0/15.0,
                        ge-0/0/16.0, ge-0/0/17.0, ge-0/0/18.0, ge-0/0/19.0*,
                        ge-0/0/20.0, ge-0/0/21.0, ge-0/0/22.0, ge-0/0/23.0,
                        ge-0/1/0.0*, ge-0/1/1.0*, ge-0/1/2.0*, ge-0/1/3.0*
mgmt
                        me0.0*
```

Meaning The **show vlans** command lists the VLANs configured on the switch. This output shows that the VLAN **default** has been created.

Verifying That Interfaces Are Associated with the Proper VLANs

Purpose Verify that Ethernet switching is enabled on switch interfaces and that all interfaces are included in the VLAN.

Action List all interfaces on which switching is enabled:

```
user@switch> show ethernet-switching interfaces

Interface   State    VLAN members           Blocking
ge-0/0/0.0  up       default                unblocked
ge-0/0/1.0  down     default                blocked - blocked by STP/RTG
ge-0/0/2.0  down     default                blocked - blocked by STP/RTG
ge-0/0/3.0  down     default                blocked - blocked by STP/RTG
ge-0/0/4.0  down     default                blocked - blocked by STP/RTG
```

```
ge-0/0/5.0  down     default              blocked - blocked by STP/RTG
ge-0/0/6.0  down     default              blocked - blocked by STP/RTG
ge-0/0/7.0  down     default              blocked - blocked by STP/RTG
ge-0/0/8.0  up       default              unblocked
ge-0/0/9.0  down     default              blocked - blocked by STP/RTG
ge-0/0/10.0 down     default              blocked - blocked by STP/RTG
ge-0/0/11.0 up       default              unblocked
ge-0/0/12.0 down     default              blocked - blocked by STP/RTG
ge-0/0/13.0 down     default              blocked - blocked by STP/RTG
ge-0/0/14.0 down     default              blocked - blocked by STP/RTG
ge-0/0/15.0 down     default              blocked - blocked by STP/RTG
ge-0/0/16.0 down     default              blocked - blocked by STP/RTG
ge-0/0/17.0 down     default              blocked - blocked by STP/RTG
ge-0/0/18.0 down     default              blocked - blocked by STP/RTG
ge-0/0/19.0 up       default              unblocked
ge-0/0/20.0 down     default              blocked - blocked by STP/RTG
ge-0/0/21.0 down     default              blocked - blocked by STP/RTG
ge-0/0/22.0 down     default              blocked - blocked by STP/RTG
ge-0/0/23.0 down     default              blocked - blocked by STP/RTG
ge-0/1/0.0  up       default              unblocked
ge-0/1/1.0  up       default              unblocked
ge-0/1/2.0  up       default              unblocked
ge-0/1/3.0  up       default              unblocked
me0.0       up       mgmt                 unblocked
```

Meaning The **show ethernet-switching interfaces** command lists all interfaces on which switching is enabled (in the **Interfaces** column), along with the VLANs that are active on the interfaces (in the **VLAN members** column). The output in this example shows all the connected interfaces, **ge-0/0/0** through **ge-0/0/12** and **ge-0/0/17** through **ge-0/0/20** and that they are all part of VLAN **default**. Notice that the interfaces listed are the logical interfaces, not the physical interfaces. For example, the output shows **ge-0/0/0.0** instead of **ge-0/0/0**. This is because the Junos OS creates VLANs on logical interfaces, not directly on physical interfaces.

Related Documentation

- Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches on page 1070
- Example: Connecting an Access Switch to a Distribution Switch on page 1078
- Example: Configure Automatic VLAN Administration Using GVRP on page 1087
- Understanding Bridging and VLANs on J-EX Series Switches on page 1041

Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches

To segment traffic on a LAN into separate broadcast domains, you create separate virtual LANs (VLANs) on a J-EX Series switch. Each VLAN is a collection of network nodes. When you use VLANs, frames whose origin and destination are in the same VLAN are forwarded only within the local VLAN, and only frames not destined for the local VLAN are forwarded to other broadcast domains. VLANs thus limit the amount of traffic flowing across the entire LAN, reducing the possible number of collisions and packet retransmissions within the LAN.

This example describes how to configure bridging for a J-EX Series switch and how to create two VLANs to segment the LAN:

Requirements

This example uses the following hardware and software components:

- One J-EX4200-48T Virtual Chassis switch

Before you set up bridging and VLANs, be sure you have:

- Installed the J-EX Series switch. See Installing and Connecting a J-EX4200 Switch.
- Performed the initial switch configuration. See “Connecting and Configuring a J-EX Series Switch (J-Web Procedure)” on page 163.

Overview and Topology

J-EX Series switches connect all devices in an office or data center into a single LAN to provide sharing of common resources such as printers and file servers and to enable wireless devices to connect to the LAN through wireless access points. The default configuration creates a single VLAN, and all traffic on the switch is part of that broadcast domain. Creating separate network segments reduces the span of the broadcast domain and allows you to group related users and network resources without being limited by physical cabling or by the location of a network device in the building or on the LAN.

This example shows a simple configuration to illustrate the basic steps for creating two VLANs on a single switch. One VLAN, called **sales**, is for the sales and marketing group, and a second, called **support**, is for the customer support team. The sales and support groups each have their own dedicated file servers, printers, and wireless access points. For the switch ports to be segmented across the two VLANs, each VLAN must have its own broadcast domain, identified by a unique name and tag (VLAN ID). In addition, each VLAN must be on its own distinct IP subnet.

The topology for this example consists of one J-EX4200-48T switch, which has a total of 48 Gigabit Ethernet ports, 8 which support Power over Ethernet (PoE). Some of the switch ports connect to Avaya IP telephones. Other ports connect to wireless access points, file servers, and printers.

Table 143: Components of the Multiple VLAN Topology

Property	Settings
Switch hardware	J-EX4200-48T, 48 Gigabit Ethernet ports, 8 of them PoE-enabled (**ge-0/0/0** through **ge-0/0/07**)

Table 143: Components of the Multiple VLAN Topology *(continued)*

Property	Settings
VLAN names and tag IDs	**sales**, tag **100** **support**, tag **200**
VLAN subnets	**sales**: **192.0.2.0/25** (addresses **192.0.2.1** through **192.0.2.126**) **support**: **192.0.2.128/25** (addresses **192.0.2.129** through **192.0.2.254**)
Interfaces in VLAN **sales**	Avaya IP telephones: **ge-0/0/2** through **ge-0/0/4** Wireless access point: **ge-0/0/0** Printers: **ge-0/0/22** and **ge-0/0/23** File servers: **ge-0/0/20** and **ge-0/0/21**
Interfaces in VLAN **support**	Avaya IP telephones: **ge-0/0/5** through **ge-0/0/7** Wireless access point: **ge-0/0/1** Printers: **ge-0/0/44** and **ge-0/0/45** File servers: **ge-0/0/46** and **ge-0/0/47**
Unused interfaces	**ge-0/0/8** through **ge0/0/19** and **ge-0/0/24** through **ge–0/0/43**

This configuration example creates two IP subnets, one for the sales VLAN and the second for the support VLAN. The switch bridges traffic within a VLAN. For traffic passing between two VLANs, the switch routes the traffic using a Layer 3 routing interface on which you have configured the address of the IP subnet.

To keep the example simple, the configuration steps show only a few devices in each of the VLANs. Use the same configuration procedure to add more LAN devices.

Configuration

Configure Layer 2 switching for two VLANs:

CLI Quick Configuration

To quickly configure Layer 2 switching for the two VLANs (**sales** and **support**) and to quickly configure Layer 3 routing of traffic between the two VLANs, copy the following commands and paste them into the switch terminal window:

```
[edit]
set interfaces ge-0/0/0 unit 0 description “Sales wireless access point port”
set interfaces ge-0/0/0 unit 0 family ethernet-switching vlan members sales
set interfaces ge-0/0/3 unit 0 description “Sales phone port”
set interfaces ge-0/0/3 unit 0 family ethernet-switching vlan members sales
set interfaces ge-0/0/22 unit 0 description “Sales printer port”
set interfaces ge-0/0/22 unit 0 family ethernet-switching vlan members sales
set interfaces ge-0/0/20 unit 0 description “Sales file server port”
set interfaces ge-0/0/20 unit 0 family ethernet-switching vlan members sales
set interfaces ge-0/0/1 unit 0 description “Support wireless access point port”
set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan members support
set interfaces ge-0/0/6 unit 0 description “Support phone port”
set interfaces ge-0/0/6 unit 0 family ethernet-switching vlan members support
set interfaces ge-0/0/44 unit 0 description “Support printer port”
set interfaces ge-0/0/44 unit 0 family ethernet-switching vlan members support
set interfaces ge-0/0/46 unit 0 description “Support file server port”
set interfaces ge-0/0/46 unit 0 family ethernet-switching vlan members support
```

```
set interfaces vlan unit 0 family inet address 192.0.2.0/25
set interfaces vlan unit 1 family inet address 192.0.2.128/25
set vlans sales l3–interface vlan.0
set vlans sales vlan-id 100
set vlans support vlan-id 200
set vlans support l3-interface vlan.1
```

Step-by-Step Procedure

Configure the switch interfaces and the VLANs to which they belong. By default, all interfaces are in access mode, so you do not have to configure the port mode.

1. Configure the interface for the wireless access point in the sales VLAN:

```
[edit interfaces ge-0/0/0 unit 0]
user@switch# set description “Sales wireless access point port”
user@switch# set family ethernet-switching vlan members sales
```

2. Configure the interface for the Avaya IP phone in the sales VLAN:

```
[edit interfaces ge-0/0/3 unit 0]
user@switch# set description “Sales phone port”
user@switch# set family ethernet-switching vlan members sales
```

3. Configure the interface for the printer in the sales VLAN:

```
[edit interfaces ge-0/0/22 unit 0]
user@switch# set description “Sales printer port”
user@switch# set family ethernet-switching vlan members sales
```

4. Configure the interface for the file server in the sales VLAN:

```
[edit interfaces ge-0/0/20 unit 0]
user@switch# set description “Sales file server port”
user@switch# set family ethernet-switching vlan members sales
```

5. Configure the interface for the wireless access point in the support VLAN:

```
[edit interfaces ge-0/0/1 unit 0]
user@switch# set description “Support wireless access point port”
user@switch# set family ethernet-switching vlan members support
```

6. Configure the interface for the Avaya IP phone in the support VLAN:

```
[edit interfaces ge-0/0/6 unit 0]
user@switch# set description “Support phone port”
user@switch# set family ethernet-switching vlan members support
```

7. Configure the interface for the printer in the support VLAN:

```
[edit interfaces ge-0/0/44 unit 0]
user@switch# set description “Support printer port”
user@switch# set family ethernet-switching vlan members support
```

8. Configure the interface for the file server in the support VLAN:

```
[edit interfaces ge-0/0/46 unit 0]
user@switch# set description “Support file server port”
user@switch# set family ethernet-switching vlan members support
```

9. Create the subnet for the sales broadcast domain:

```
[edit interfaces]
user@switch# set vlan unit 0 family inet address 192.0.2.1/25
```

10. Create the subnet for the support broadcast domain:

```
[edit interfaces]
user@switch# set vlan unit 1 family inet address 192.0.2.129/25
```

11. Configure the VLAN tag IDs for the sales and support VLANs:

```
[edit vlans]
user@switch# set sales vlan-id 100
user@switch# set support vlan-id 200
```

12. To route traffic between the sales and support VLANs, define the interfaces that are members of each VLAN and associate a Layer 3 interface:

```
[edit vlans]
user@switch# set sales l3-interface
user@switch# set support l3-interface vlan.1
```

Display the results of the configuration:

```
user@switch> show configuration
interfaces {
  ge-0/0/0 {
    unit 0 {
      description “Sales wireless access point port”;
      family ethernet-switching {
        vlan members sales;
      }
    }
  }
  ge-0/0/3 {
    unit 0 {
      description “Sales phone port”;
      family ethernet-switching {
        vlan members sales;
      }
    }
  }
  ge-0/0/22 {
    unit 0 {
      description “Sales printer port”;
      family ethernet-switching {
        vlan members sales;
      }
    }
  }
  ge-0/0/20 {
    unit 0 {
      description “Sales file server port”;
      family ethernet-switching {
        vlan members sales;
      }
    }
  }
  ge-0/0/1 {
    unit 0 {
      description “Support wireless access point port”;
```

```
        family ethernet-switching {
          vlan members support;
        }
      }
    }
    ge-0/0/6 {
      unit 0 {
        description “Support phone port”;
        family ethernet-switching {
          vlan members support;
        }
      }
    }
    ge-0/0/44 {
      unit 0 {
        description “Support printer port”;
        family ethernet-switching {
          vlan members support;
        }
      }
    }
    ge-0/0/46 {
      unit 0 {
        description “Support file server port”;
        family ethernet-switching {
          vlan members support;
        }
      }
      vlans {
        unit 0 {
          family inet address 192.0.2.0/25;
        }
        unit 1 {
          family inet address 192.0.2.128/25;
        }
      }
    }
}
vlans {
  sales {
    vlan-id 100;
    interface ge-0/0/0.0:
    interface ge-0/0/3/0;
    interface ge-0/0/20.0;
    interface ge-0/0/22.0;
    l3-interface vlan 0;
  }
  support {
    vlan-id 200;
    interface ge-0/0/1..0:
    interface ge-0/0/6.0;
    interface ge-0/0/44.0;
    interface ge-0/0/46.0;
    l3-interface vlan 1;
  }
}
```

TIP: To quickly configure the sales and support VLAN interfaces, issue the load merge terminal command, then copy the hierarchy and paste it into the switch terminal window.

Verification

To verify that the “sales” and “support” VLANs have been created and are operating properly, perform these tasks:

- Verifying That the VLANs Have Been Created and Associated to the Correct Interfaces on page 1076
- Verifying That Traffic Is Being Routed Between the Two VLANs on page 1077
- Verifying That Traffic Is Being Switched Between the Two VLANs on page 1077

Verifying That the VLANs Have Been Created and Associated to the Correct Interfaces

Purpose Verify that the VLANs **sales** and **support** have been created on the switch and that all connected interfaces on the switch are members of the correct VLAN.

Action List all VLANs configured on the switch:

Use the operational mode commands:

```
user@switch>  show vlans
Name            Tag     Interfaces
default
                        ge-0/0/1.0, ge-0/0/2.0, ge-0/0/4.0, ge-0/0/5.0,
                        ge-0/0/6.0,  ge-0/0/7.0, ge-0/0/8.0, ge-0/0/9.0,
                        ge-0/0/10.0*, ge-0/0/11.0, ge-0/0/12.0, ge-0/0/13.0*,
                        ge-0/0/14.0, ge-0/0/15.0, ge-0/0/16.0, ge-0/0/17.0,
                        ge-0/0/18.0, ge-0/0/19.0, ge-0/0/21.0, ge-0/0/23.0*,
                        ge-0/0/25.0, ge-0/0/27.0, ge-0/0/28.0, ge-0/0/29.0,
                        ge-0/0/30.0, ge-0/0/31.0, ge-0/0/32.0, ge-0/0/33.0,
                        ge-0/0/34.0, ge-0/0/35.0, ge-0/0/36.0, ge-0/0/37.0,
                        ge-0/0/38.0, ge-0/0/39.0, ge-0/0/40.0, ge-0/0/41.0,
                        ge-0/0/42.0, ge-0/0/43.0, ge-0/0/45.0, ge-0/0/47.0,
                        ge-0/1/0.0*, ge-0/1/1.0*, ge-0/1/2.0*, ge-0/1/3.0*

sales       100
                        ge-0/0/0.0*, ge-0/0/3.0, ge-0/0/20.0, ge-0/0/22.0

support     200
                        ge-0/0/1.0, ge-0/0/6.0, ge-0/0/44.0, ge-0/0/46.0*
mgmt
                        me0.0*
```

Meaning The **show vlans** command lists all VLANs configured on the switch and which interfaces are members of each VLAN. This command output shows that the **sales** and **support** VLANs have been created. The **sales** VLAN has a tag ID of 100 and is associated with interfaces **ge-0/0/0.0**, **ge-0/0/3.0**, **ge-0/0/20.0**, and **ge-0/0/22.0**. VLAN **support** has a

tag ID of 200 and is associated with interfaces **ge-0/0/1.0**, **ge-0/0/6.0**, **ge-0/0/44.0**, and **ge-0/0/46.0**.

Verifying That Traffic Is Being Routed Between the Two VLANs

Purpose Verify routing between the two VLANs.

Action List the Layer 3 routes in the switch's Address Resolution Protocol (ARP) table:

```
user@switch> show arp
MAC Address       Address         Name                      Flags

00:00:0c:06:2c:0d       192.0.2.3       vlan.0                    None
00:13:e2:50:62:e0       192.0.2.11      vlan.1                    None
```

Meaning Sending IP packets on a multiaccess network requires mapping from an IP address to a MAC address (the physical or hardware address). The ARP table displays the mapping between the IP address and MAC address for both **vlan.0** (associated with **sales**) and **vlan.1** (associated with **support**). These VLANs can route traffic to each other.

Verifying That Traffic Is Being Switched Between the Two VLANs

Purpose Verify that learned entries are being added to the Ethernet switching table.

Action List the contents of the Ethernet switching table:

```
user@switch> show ethernet-switching table

Ethernet-switching table: 8 entries, 5 learned
  VLAN              MAC address       Type         Age Interfaces
  default           *                 Flood          - All-members
  default           00:00:05:00:00:01 Learn          - ge-0/0/10.0
  default           00:00:5e:00:01:09 Learn          - ge-0/0/13.0
  default           00:19:e2:50:63:e0 Learn          - ge-0/0/23.0
  sales             *                 Flood          - All-members
  sales             00:00:5e:00:07:09 Learn          - ge-0/0/0.0
  support           *                 Flood          - All-members
  support           00:00:5e:00:01:01 Learn          - ge-0/0/46.0
```

Meaning The output shows that learned entries for the **sales** and **support** VLANs have been added to the Ethernet switching table, and are associated with interfaces **ge-0/0/0.0** and **ge-0/0/46.0**. Even though the VLANs were associated with more than one interface in the configuration, these interfaces are the only ones that are currently operating.

Related Documentation

- Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch on page 1063
- Example: Connecting an Access Switch to a Distribution Switch on page 1078
- Example: Configure Automatic VLAN Administration Using GVRP on page 1087
- Understanding Bridging and VLANs on J-EX Series Switches on page 1041

Example: Connecting an Access Switch to a Distribution Switch

In large local area networks (LANs), you commonly need to aggregate traffic from a number of access switches into a distribution switch.

This example describes how to connect an access switch to a distribution switch:

Requirements

This example uses the following hardware and software components:

- For the distribution switch, one J-EX4200-24F switch. This model is designed to be used as a distribution switch for aggregation or collapsed core network topologies and in space-constrained data centers. It has twenty-four 1-Gigabit Ethernet fiber SFP ports and an uplink module with two 10-Gigabit Ethernet ports.
- For the access switch, one J-EX4200-24T, which has twenty-four 1-Gigabit Ethernet ports, 8 of which support Power over Ethernet (PoE), and an uplink module with four 1-Gigabit Ethernet ports.

Before you connect an access switch to a distribution switch, be sure you have:

- Installed the two switches. See Installing and Connecting a J-EX4200 Switch.
- Performed the initial software configuration on both switches. See “Connecting and Configuring a J-EX Series Switch (J-Web Procedure)” on page 163.

Overview and Topology

In a large office that is spread across several floors or buildings, or in a data center, you commonly aggregate traffic from a number of access switches into a distribution switch. This configuration example shows a simple topology to illustrate how to connect a single access switch to a distribution switch.

In the topology, the LAN is segmented into two VLANs, one for the sales department and the second for the support team. One 1-Gigabit Ethernet port on the access switch's uplink module connects to the distribution switch, to one 1-Gigabit Ethernet port on the distribution switch.

Figure 28 on page 1079 shows one J-EX4200 switch that is connected to the three access switches.

Figure 28: Topology for Configuration

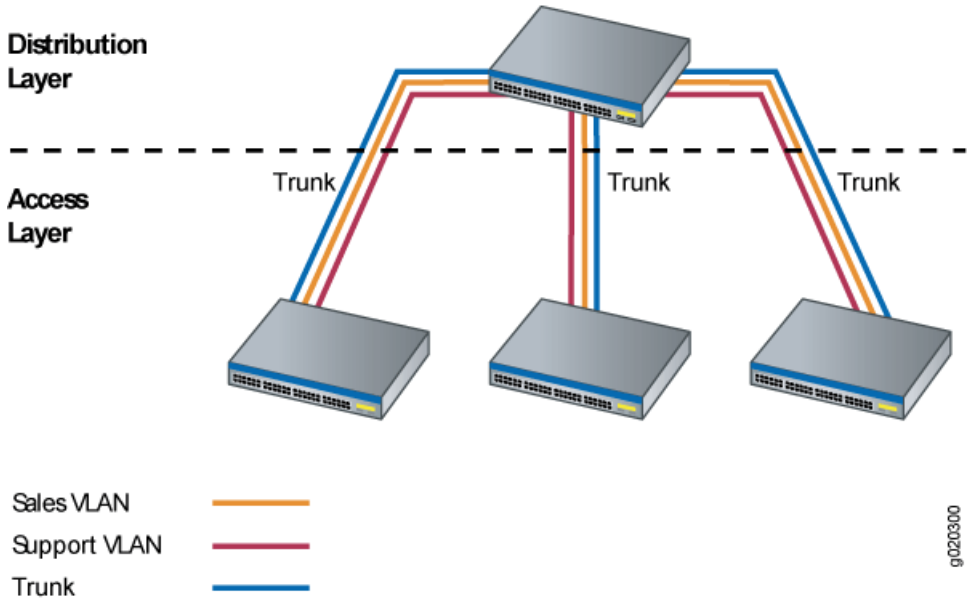

Table 144 on page 1079 explains the components of the example topology. The example shows how to configure one of the three access switches. The other access switches could be configured in the same manner.

Table 144: Components of the Topology for Connecting an Access Switch to a Distribution Switch

Property	Settings
Access switch hardware	J-EX4200-24T, 24 1-Gigabit Ethernet ports, with 8 ports PoE-enabled (**ge-0/0/0** through **ge-0/0/7**); one uplink module
Distribution switch hardware	J-EX4200-24F, 24 1-Gigabit Ethernet fiber SPF ports (**ge-0/0/0** through **ge-0/0/23**); one uplink module
VLAN names and tag IDs	**sales**, tag **100** **support**, tag **200**
VLAN subnets	**sales**: **192.0.2.0/25** (addresses **192.0.2.1** through **192.0.2.126**) **support**: **192.0.2.128/25** (addresses **192.0.2.129** through **192.0.2.254**)
Trunk port interfaces	On the access switch: **ge-0/1/0** On the distribution switch: **ge-0/0/0**
Access port interfaces in VLAN **sales** (on access switch)	Avaya IP telephones: **ge-0/0/2** through **ge-0/0/4** Wireless access point: **ge-0/0/0** Printers: **ge-0/0/22** and **ge-0/0/23** File servers: **ge-0/0/20** and **ge-0/0/21**

Table 144: Components of the Topology for Connecting an Access Switch to a Distribution Switch *(continued)*

Property	Settings
Access port interfaces in VLAN **support** (on access switch)	Avaya IP telephones: **ge-0/0/5** through **ge-0/0/7** Wireless access point: **ge-0/0/1** Printers: **ge-0/0/44** and **ge-0/0/45** File servers: **ge-0/0/46** and **ge-0/0/47**
Unused interfaces on access switch	**ge-0/0/8** through **ge-0/0/19** and **ge-0/0/24** through **ge-0/0/43**

Configuring the Access Switch

To configure the access switch:

CLI Quick Configuration

To quickly configure the access switch, copy the following commands and paste them into the switch terminal window:

```
[edit]
set interfaces ge-0/0/0 unit 0 description “Sales Wireless access point port”
set interfaces ge-0/0/0 unit 0 family ethernet-switching vlan members sales
set interfaces ge-0/0/3 unit 0 description “Sales phone port”
set interfaces ge-0/0/3 unit 0 family ethernet-switching vlan members sales
set interfaces ge-0/0/22 unit 0 description “Sales printer port”
set interfaces ge-0/0/22 unit 0 family ethernet-switching vlan members sales
set interfaces ge-0/0/20 unit 0 description “Sales file server port”
set interfaces ge-0/0/20 unit 0 family ethernet-switching vlan members sales
set interfaces ge-0/0/1 unit 0 description “Support wireless access point port”
set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan members support
set interfaces ge-0/0/6 unit 0 description “Support phone port”
set interfaces ge-0/0/6 unit 0 family ethernet-switching vlan members support
set interfaces ge-0/0/44 unit 0 description “Support printer port”
set interfaces ge-0/0/44 unit 0 family ethernet-switching vlan members support
set interfaces ge-0/0/46 unit 0 description “Support file server port”
set interfaces ge-0/0/46 unit 0 family ethernet-switching vlan members support
set interfaces ge-0/1/0 unit 0 description “Uplink module port connection to distribution switch”
set interfaces ge-0/1/0 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/1/0 unit 0 family ethernet-switching native-vlan-id 1
set interfaces ge-0/1/0 unit 0 family ethernet switching vlan members [sales support]
set interfaces vlan unit 0 family inet address 192.0.2.1/25
set interfaces vlan unit 1 family inet address 192.0.2.129/25
set vlans sales interface ge-0/0/0.0
set vlans sales interface ge-0/0/3.0
set vlans sales interface ge-0/0/22.0
set vlans sales interface ge-0/0/20.0
set vlans sales l3-interface vlan.0
set vlans sales vlan-id 100
set vlans sales vlan-description “Sales VLAN”
set vlans support interface ge-0/0/1.0
set vlans support interface ge-0/0/6.0
set vlans support interface ge-0/0/44.0
set vlans support interface ge-0/0/46.0
set vlans support vlan-id 200
set vlans support l3–interface vlan.1
set vlans support vlan-description “Support VLAN”
```

Step-by-Step Procedure

To configure the access switch:

1. Configure the 1-Gigabit Ethernet interface on the uplink module to be the trunk port that connects to the distribution switch:

```
[edit interfaces ge-0/1/0 unit 0]
user@access-switch# set description “Uplink module port connection to distribution switch”
user@access-switch# set ethernet-switching port-mode trunk
```

2. Specify the VLANs to be aggregated on the trunk port:

```
[edit interfaces ge-0/1/0 unit 0]
user@access-switch# set ethernet-switching vlan members [ sales support ]
```

3. Configure the VLAN ID to use for packets that are received with no dot1q tag (untagged packets):

```
[edit interfaces ge-0/1/0 unit 0]
user@access-switch# set ethernet-switching native-vlan-id 1
```

4. Configure the sales VLAN:

```
[edit vlans sales]
user@access-switch# set vlan-description “Sales VLAN”
user@access-switch# set vlan-id 100
user@access-switch# set l3-interface vlan.0
```

5. Configure the support VLAN:

```
[edit vlans support]
user@access-switch# set vlan-description “Support VLAN”
user@access-switch# set vlan-id 200
user@access-switch# set l3-interface vlan.1
```

6. Create the subnet for the sales broadcast domain:

```
[edit interfaces]
user@access-switch# set vlan unit 0 family inet address 192.0.2.1/25
```

7. Create the subnet for the support broadcast domain:

```
[edit interfaces]
user@access-switch# set vlan unit 1 family inet address 192.0.2.129/25
```

8. Configure the interfaces in the sales VLAN:

```
[edit interfaces]
user@access-switch# set ge-0/0/0 unit 0 description “Sales wireless access point port”
user@access-switch# set ge-0/0/0 unit 0 family ethernet-switching vlan members sales
user@access-switch# set ge-0/0/3 unit 0 description “Sales phone port”
user@access-switch# set ge-0/0/3 unit 0 family ethernet-switching vlan members sales
user@access-switch# set ge-0/0/20 unit 0 description “Sales file server port”
user@access–switch# set ge-0/0/20 unit 0 family ethernet-switching vlan members sales
user@access-switch# set ge-0/0/22 unit 0 description “Sales printer port”
user@access-switch# set ge-0/0/22 unit 0 family ethernet-switching vlan members sales
```

9. Configure the interfaces in the support VLAN:

```
[edit interfaces]
user@access-switch# set ge-0/0/1 unit 0 description “Support wireless access point
port”
user@access-switch# set ge-0/0/1 unit 0 family ethernet-switching vlan members
support
user@access-switch# set ge-0/0/6 unit 0 description “Support phone port”
user@access-switch# set ge-0/0/6 unit 0 family ethernet-switching vlan members
support
user@access-switch# set ge-0/0/44 unit 0 description “Support printer port”
user@access-switch# set ge-0/0/44 unit 0 family ethernet-switching vlan members
support
user@access-switch# set ge-0/0/46 unit 0 description “Support file server port”
user@access-switch# set ge-0/0/46 unit 0 family ethernet-switching vlan members
support
```

10. Configure descriptions and VLAN tag IDs for the sales and support VLANs:

```
[edit vlans]
user@access-switch# set sales vlan-description “Sales VLAN”
user@access-switch# set sales vlan-id 100
user@access-switch# set support vlan-description “Support VLAN”
user@access-switch# set support vlan-id 200
```

11. To route traffic between the sales and support VLANs and associate a Layer 3 interface with each VLAN:

```
[edit vlans]
user@access-switch# set sales l3-interface vlan.0
user@access-switch# set support l3-interface vlan.1
```

Results Display the results of the configuration:

```
user@access-switch> show
interfaces {
  ge-0/0/0 {
    unit 0 {
      description “Sales wireless access point port”;
      family ethernet-switching {
        vlan members sales;
      }
    }
  }
  ge-0/0/3 {
    unit 0 {
      description “Sales phone port”;
        family ethernet-switching {
        vlan members sales;
      }
    }
  }
  ge-0/0/20 {
    unit 0 {
      description “Sales file server port”;
        family ethernet-switching {
        vlan members sales;
      }
```

```
    }
  }
  ge-0/0/22 {
    unit 0 {
      description “Sales printer port”;
        family ethernet-switching {
        vlan members sales;
      }
    }
  }
  ge-0/0/1 {
    unit 0 {
      description “Support wireless access point port”;
        family ethernet-switching {
        vlan members support;
      }
    }
  }
  ge-0/0/6 {
    unit 0 {
      description “Support phone port”;
        family ethernet-switching {
        vlan members support;
      }
    }
  }
  ge-0/0/44 {
    unit 0 {
      description “Support printer port”;
        family ethernet-switching {
        vlan members sales;
      }
    }
  }
  ge-0/0/46 {
    unit 0 {
      description “Support file server port”;
        family ethernet-switching {
        vlan members support;
      }
    }
  }
  ge-0/1/0 {
    unit 0 {
      description “Uplink module port connection to distribution switch”;
        family ethernet-switching {
        port-mode trunk;
          vlan members [ sales support ];
        native-vlan-id 1;
      }
    }
  }
  vlan {
    unit 0 {
      family inet address 192.0.2.1/25;
    }
```

```
        unit 1 {
          family inet address 192.0.2.129/25;
        }
      }
    }
    vlans {
      sales {
        vlan-id 100;
        vlan-description “Sales VLAN”;
        l3-interface vlan.0;
      }
      support {
        vlan-id 200;
        vlan-description “Support VLAN”;
        l3-interface vlan.1;
      }
    }
```

TIP: To quickly configure the distribution switch, issue the load merge terminal command, then copy the hierarchy and paste it into the switch terminal window.

Configuring the Distribution Switch

To configure the distribution switch:

CLI Quick Configuration

To quickly configure the distribution switch, copy the following commands and paste them into the switch terminal window:

```
set interfaces ge-0/0/0 description “Connection to access switch”
set interfaces ge-0/0/0 ethernet-switching port-mode trunk
set interfaces ge-0/0/0 ethernet-switching vlan members [ sales support ]
set interfaces ge-0/0/0 ethernet-switching native-vlan-id 1
set interfaces vlan unit 0 family inet address 192.0.2.2/25
set interfaces vlan unit 1 family inet address 192.0.2.130/25
set vlans sales vlan-description “Sales VLAN”
set vlans sales vlan-id 100
set vlans sales l3-interface vlan.0
set vlans support vlan-description “Support VLAN”
set vlans support vlan-id 200
set vlans support l3-interface vlan.1
```

Step-by-Step Procedure

To configure the distribution switch:

1. Configure the interface on the switch to be the trunk port that connects to the access switch:

   ```
   [edit interfaces ge-0/0/0 unit 0]
   user@distribution-switch# set description “Connection to access switch”
   user@distribution-switch# set ethernet-switching port-mode trunk
   ```

2. Specify the VLANs to be aggregated on the trunk port:

   ```
   [edit interfaces ge-0/0/0 unit 0]
   user@distribution-switch# set ethernet-switching vlan members [ sales support ]
   ```

3. Configure the VLAN ID to use for packets that are received with no dot1q tag (untagged packets):

```
[edit interfaces]
user@distribution-switch# set ge-0/0/0 ethernet-switching native-vlan-id 1
```

4. Configure the sales VLAN:

```
[edit vlans sales]
user@distribution-switch# set vlan-description “Sales VLAN”
user@distribution-switch# set vlan-id 100
user@distribution-switch# set l3-interface vlan.0
```

5. Configure the support VLAN:

```
[edit vlans support]
user@distribution-switch# set vlan-description “Support VLAN”
user@distribution-switch# set vlan-id 200
user@distribution-switch# set l3-interface vlan.1
```

6. Create the subnet for the sales broadcast domain:

```
[edit interfaces]
user@distribution-switch# set vlan unit 0 family inet address 192.0.2.2/25
```

7. Create the subnet for the support broadcast domain:

```
[edit interfaces]
user@distribution-switch# set vlan unit 1 family inet address 192.0.2.130/25
```

Results Display the results of the configuration:

```
user@distribution-switch> show
interfaces {
    ge-0/0/0 {
        description “Connection to access switch”;
        unit 0 {
            family ethernet-switching {
                port-mode trunk;
                vlan members [ sales support ];
                native-vlan-id 1;
            }
        }
    }
    vlan {
        unit 0 {
            family inet address 192.0.2.2/25;
        }
        unit 1 {
            family inet address 192.0.2.130/25;
        }
    }
}
vlans {
    sales {
        vlan-id 100;
        vlan-description “Sales VLAN”;
        l3-interface vlan.0;
    }
```

```
    support {
      vlan-id 200;
      vlan-description “Support VLAN”;
      l3-interface vlan.1;
    }
  }
```

TIP: To quickly configure the distribution switch, issue the load merge terminal command, then copy the hierarchy and paste it into the switch terminal window.

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying the VLAN Members and Interfaces on the Access Switch on page 1086
- Verifying the VLAN Members and Interfaces on the Distribution Switch on page 1087

Verifying the VLAN Members and Interfaces on the Access Switch

Purpose Verify that the **sales** and **support** have been created on the switch.

Action List all VLANs configured on the switch:

```
user@switch> show vlans

Name            Tag     Interfaces
default
                        ge-0/0/1.0, ge-0/0/2.0, ge-0/0/4.0, ge-0/0/5.0,
                        ge-0/0/6.0, ge-0/0/7.0, ge-0/0/8.0*, ge-0/0/9.0,

                        ge-0/0/10.0, ge-0/0/11.0*, ge-0/0/12.0, ge-0/0/13.0,
                        ge-0/0/14.0, ge-0/0/15.0, ge-0/0/16.0, ge-0/0/17.0,
                        ge-0/0/18.0, ge-0/0/19.0*,ge-0/0/21.0, ge-0/0/23.0,
                        ge-0/0/25.0, ge-0/0/27.0*,ge-0/0/28.0, ge-0/0/29.0,
                        ge-0/0/30.0, ge-0/0/31.0*,ge-0/0/32.0, ge-0/0/33.0,
                        ge-0/0/34.0, ge-0/0/35.0*,ge-0/0/36.0, ge-0/0/37.0,
                        ge-0/0/38.0, ge-0/0/39.0*,ge-0/0/40.0, ge-0/0/41.0,
                        ge-0/0/42.0, ge-0/0/43.0*,ge-0/0/45.0, ge-0/0/47.0,
                        ge-0/1/1.0*, ge-0/1/2.0*, ge-0/1/3.0*

sales           100
                        ge-0/0/0.0*, ge-0/0/3.0, ge-0/0/20.0, ge-0/0/22.0,
                        ge-0/1/0.0*,

support         200
                        ge-0/0/1.0*, ge-0/0/6.0, ge-0/0/44.0, ge-0/0/46.0,

mgmt
                        me0.0*
```

Meaning The output shows the **sales** and **support** VLANs and the interfaces associated with them.

Verifying the VLAN Members and Interfaces on the Distribution Switch

Purpose Verify that the **sales** and **support** have been created on the switch.

Action List all VLANs configured on the switch:

```
user@switch> show vlans

Name            Tag     Interfaces
default
                        ge-0/0/1.0, ge-0/0/2.0, ge-0/0/3.0, ge-0/0/4.0,
                        ge-0/0/5.0, ge-0/0/6.0, ge-0/0/7.0*, ge-0/0/8.0,

                        ge-0/0/9.0, ge-0/0/10.0*, ge-0/0/11.0, ge-0/0/12.0,
                        ge-0/0/13.0, ge-0/0/14.0, ge-0/0/15.0, ge-0/0/16.0,
                        ge-0/0/17.0, ge-0/0/18.0*, ge-0/0/19.0, ge-0/0/20.0,
                        ge-0/0/21.0, ge-0/0/22.0*, ge-0/0/23.0, ge-0/1/1.0*,
                        ge-0/1/2.0*, ge-0/1/3.0*

sales           100
                        ge-0/0/0.0*

support         200
                        ge-0/0/0.0*

mgmt
                        me0.0*
```

Meaning The output shows the **sales** and **support** VLANs associated to interface **ge-0/0/0.0**. Interface **ge-0/0/0.0** is the trunk interface connected to the access switch.

Related Documentation

- Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch on page 1063
- Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches on page 1070
- Example: Configure Automatic VLAN Administration Using GVRP on page 1087
- Understanding Bridging and VLANs on J-EX Series Switches on page 1041

Example: Configure Automatic VLAN Administration Using GVRP

As a network expands and the number of clients and VLANs increases, VLAN administration becomes complex, and the task of efficiently configuring VLANs on multiple J-EX Series switches becomes increasingly difficult. To automate VLAN administration, you can enable GARP VLAN Registration Protocol (GVRP) on the network.

NOTE: Only trunk interfaces can be enabled for GVRP.

This example describes how to use GVRP to automate administration of VLAN membership changes within your network:

Requirements

This example uses the following hardware and software components:

- Two J-EX4200 access switches
- One J-EX4200 distribution switch

Before you configure GVRP on the access switches and on the distribution switch, be sure you have:

- Performed the initial software configuration on the switches. See “Connecting and Configuring a J-EX Series Switch (J-Web Procedure)” on page 163.
- Configured the VLANs on both the access switches and on the distribution switch. (Dynamic VLAN configuration is not supported.)
- Configured a trunk interface on all the switches.

Overview and Topology

When you are setting up your network, you should configure all VLANs on all switches, even though some switches are not actively participating in a VLAN. Then enable GVRP on the trunk interface of each switch. GVRP ensures that the VLAN membership information on the trunk interface is updated as the switch’s access interfaces become active or inactive in the configured VLANs.

You do not need to take an extra step of explicitly binding a VLAN to the trunk interface. When GVRP is enabled, the trunk interface advertises all the VLANs that are active (bound to access interfaces) on that switch. A GVRP-enabled trunk interface does not advertise VLANs that have been configured on the switch but that are not currently bound to an access interface. Thus, GVRP provides the benefit of reducing network overhead—by limiting the scope of broadcast, unknown unicast, and multicast (BUM) traffic to interested devices only.

This example shows a network with three VLANs: finance, sales, and lab.

Access Switch A has been configured to support all three VLANS and all three VLANS are active, bound to interfaces that are connected to personal computers:

- **ge-0/0/1**—Connects PC1 as member of finance vlan, VLAN ID 100
- **ge-0/0/2**—Connects PC2 as member of lab vlan, VLAN ID 200

- **ge-0/0/3**—Connects PC3 as member of sales vlan, VLAN ID 300

Access Switch B has also been configured to support three VLANS. However, currently only two VLANs are active, bound to interfaces that are connected to personal computers:

- **ge-0/0/0**—Connects PC4 as member of finance vlan, VLAN ID 100
- **ge-0/0/1**—Connects PC5 as member of lab vlan, VLAN ID 200

The Distribution Switch is also configured to support the three VLANs (finance, lab, sales). However, the Distribution Switch does not have any access interfaces that are connecting devices as members of these VLANs. The Distribution Switch has two trunk interfaces:

- **xe-0/1/1**—Connects Distribution Switch to Access Switch A.
- **xe-0/1/0**—Connects Distribution Switch to Access Switch B.

Figure 29 on page 1089 shows GVRP configured on two access switches and one distribution switch.

Figure 29: GVRP Configured on Two Access Switches and One Distribution Switch for Automatic VLAN Administration

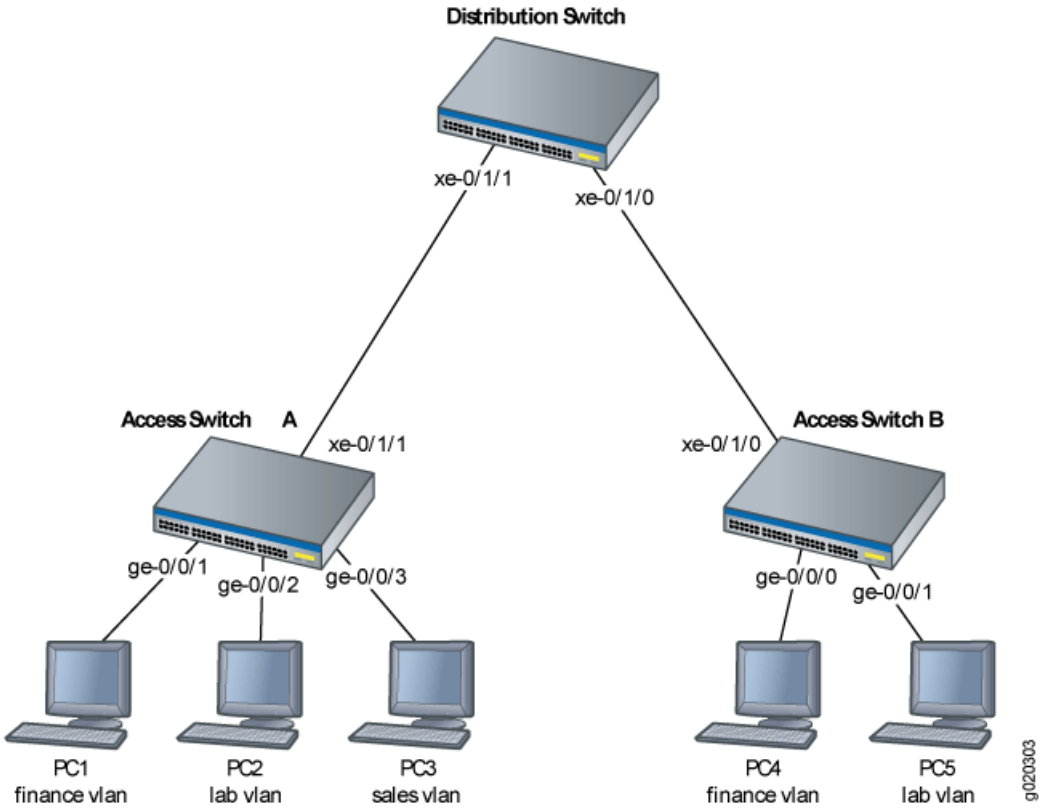

Table 145: Components of the Network Topology

Property	Settings
Switch hardware	• Access Switch A • Access Switch B • Distribution Switch
VLAN names and tag IDs	**finance**, tag **100** **lab**, tag **200** **sales**, tag **300**
Interfaces	Access Switch A Interfaces • **ge-0/0/1**—Connects PC1 to Access Switch A. • **ge-0/0/2**—Connects PC2 to Access Switch A. • **ge-0/0/3**—Connects PC3 to Access Switch A. • **xe-0/1/1**—Connects Access Switch A to Distribution Switch. (trunk) Access Switch B Interfaces • **ge-0/0/0**—Connects PC4 to Access Switch B. • **ge-0/0/1**—Connects PC5 to Access Switch B. • **xe-0/1/0**—Connects Access Switch B to Distribution Switch. (trunk) Distribution Switch Interfaces • **xe-0/1/1**—Connects Distribution Switch to Access Switch A. (trunk) • **xe-0/1/0**—Connects Distribution Switch to Access Switch B. (trunk)

When VLAN access interfaces become active or inactive, GVRP ensures that the updated information is advertised on the trunk interface. Thus, the Distribution Switch does not forward traffic to inactive VLANs.

Configuring VLANs and GVRP on Access Switch A

To configure three VLANs on the switch, bind access interfaces for PC1, PC2, and PC3 to the VLANs (finance, lab, sales), and enable GVRP on the trunk interface of Access Switch A, perform these tasks:

CLI Quick Configuration

To quickly configure Access Switch A to support the three VLANs, bind interfaces for the three PCs to the appropriate VLANs, and enable GVRP on the trunk interface, copy the following commands and paste them into the switch terminal window of Switch A:

```
[edit]
set vlans finance vlan-id 100
set vlans lab vlan-id 200
set vlans sales vlan-id 300
set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan members finance
set interfaces ge-0/0/2 unit 0 family ethernet-switching vlan members lab
set interfaces ge-0/0/3 unit 0 family ethernet-switching vlan members sales
set interfaces xe-0/1/1 unit 0 family ethernet-switching port-mode trunk
```

```
set protocols gvrp interface xe-0/1/1.0
```

NOTE: As we recommend, default GVRP timers are used in this example. The default values associated with each GVRP timer are: 200 ms for the join timer, 600 ms for the leave timer, and 1000 cs (10000 ms) for the leaveall timer. Modifying timers to inappropriate values may cause an imbalance in the operation of GVRP. Refer to IEEE 802.1D [2004] Clause 12 for more information. The timer values are displayed when you use the **show gvrp** command to verify that GVRP is enabled. For more information on the timers, see **gvrp** and its associated configuration statements.

Step-by-Step Procedure

To configure Access Switch A to support the three VLANs, bind interfaces for the three PCs to the appropriate VLANs, and enable GVRP on the trunk interface, copy the following commands and paste them into the switch terminal window of Switch A:

1. Configure the finance VLAN:

```
[edit]
user@Access-Switch-A# set vlans finance vlan-id 100
```

2. Configure the lab VLAN:

```
[edit]
user@Access-Switch-A# set vlans lab vlan–id 200
```

3. Configure the sales VLAN:

```
[edit]
user@Access-Switch-A# set vlans sales vlan–id 300
```

4. Configure an Ethernet interface as a member of the finance VLAN:

```
[edit]
does user@Access-Switch-A# set interfaces ge-0/0/1 unit 0 family ethernet-switching
vlan members finance
```

5. Configure an Ethernet interface as a member of the lab VLAN:

```
[edit]
user@Access-Switch-A# set interfaces ge-0/0/2 unit 0 family ethernet-switching vlan
members lab
```

6. Configure an Ethernet interface as a member of the sales VLAN:

```
[edit]
user@Access-Switch-A# set interfaces ge-0/0/3 unit 0 family ethernet-switching vlan
members sales
```

7. Configure a trunk interface:

```
user@Access-Switch-A# set interfaces xe-0/1/1 unit 0 family ethernet-switching
port-mode trunk
```

8. Enable GVRP on the trunk interface:

```
[edit]
user@Access-Switch-A# set protocols gvrp interface xe-0/1/1.0
```

Results Check the results of the configuration:

```
interfaces {
    ge-0/0/0 {
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/0/1 {
        unit 0 {
            family ethernet-switching {
                vlan {
                    members finance;
                }
            }
        }
    }
    ge-0/0/2 {
        unit 0 {
            family ethernet-switching {
                vlan {
                    members lab;
                }
            }
        }
    }
    ge-0/0/3 {
        unit 0 {
            family ethernet-switching {
                vlan {
                    members sales;
                }
            }
        }
    }
    xe-0/1/1 {
        unit 0 {
            family ethernet-switching {
                port-mode trunk;
            }
        }
    }
    ge-0/1/2 {
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/1/3 {
        unit 0 {
            family ethernet-switching;
        }
    }
}
    protocols {
        igmp-snooping {
            vlan all;
```

```
    }
    lldp {
      interface all;
    }
    lldp-med {
      interface all;
    }
    gvrp {
      interface xe-0/1/1.0;
    }
    rstp;
  }
  ethernet-switching-options {
    storm-control {
      interface all {
        level 50;
      }
    }
  }
  vlans {
    finance {
      vlan-id 100;
    }
    lab {
      vlan-id 200;
    }
    sales {
      vlan-id 300;
    }
```

Configuring VLANs and GVRP on Access Switch B

To configure three VLANs on the switch, bind access interfaces for PC4 and PC5 to the VLANs (finance and lab), and enable GVRP on the trunk interface of Access Switch Switch B, perform these tasks:

CLI Quick Configuration

To quickly configure Access Switch B to support the three VLANs, bind interfaces for the two PCs to the appropriate VLANs, and enable GVRP on the trunk interface, copy the following commands and paste them into the switch terminal window of Switch B:

```
[edit]
set vlans finance vlan-id 100
set vlans lab vlan-id 200
set vlans sales vlan-id 300
set interfaces ge-0/0/0 unit 0 family ethernet-switching vlan members finance
set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan members lab
set interfaces xe-0/1/0 unit 0 family ethernet-switching port-mode trunk
set protocols gvrp interface xe-0/1/0.0
```

Step-by-Step Procedure

To configure Access Switch B to support the three VLANs, bind interfaces for the two PCs to the appropriate VLAN, and enable GVRP on the trunk interface, copy the following commands and paste them into the switch terminal window of Switch B:

1. Configure the finance VLAN:

```
[edit]
user@Access-Switch-B# set vlans finance vlan-id 100
```

2. Configure the lab VLAN:

```
[edit]
user@Access-Switch-B# set vlans lab vlan–id 200
```

3. Configure the sales VLAN:

```
[edit]
user@Access-Switch-B# set vlans sales vlan–id 300
```

4. Configure an Ethernet interface as a member of the finance VLAN:

```
[edit]
user@Access-Switch-B# set interfaces ge-0/0/0 unit 0 family ethernet-switching vlan
members finance
```

5. Configure an Ethernet interface as a member of the lab VLAN:

```
[edit]
user@Access-Switch-B# set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan
members lab
```

6. Configure a trunk interface:

```
user@Access-Switch-B# set interfaces xe-0/1/0 unit 0 family ethernet-switching
port-mode trunk
```

7. Enable GVRP on the trunk interface:

```
[edit]
user@Access-Switch-B# set protocols gvrp xe-0/1/0.0
```

NOTE: As we recommend, default GVRP timers are used in this example. The default values associated with each GVRP timer are: 200 ms for the join timer, 600 ms for the leave timer, and 1000 cs (10000 ms) for the leaveall timer. Modifying timers to inappropriate values might cause an imbalance in the operation of GVRP. Refer to IEEE 802.1D [2004] Clause 12 for more information. The timer values are displayed when you use the show gvrp command to verify that GVRP is enabled. For more information on the timers, see gvrp and its associated configuration statements.

Results

Check the results of the configuration:

```
[edit]
  user@Access-Switch-B #show
interfaces {
```

```
ge-0/0/0 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members finance;
      }
    }
  }
}
ge-0/0/1 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members lab;
      }
    }
  }
}
ge-0/0/2 {
  unit 0 {
    family ethernet-switching;
  }
}
ge-0/0/3 {
  unit 0 {
    family ethernet-switching;
  }
}
ge-0/1/0 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/1/0 {
  unit 0 {
    family ethernet-switching {
      port-mode trunk;
    }
  }
}
ge-0/1/1 {
  unit 0 {
    family ethernet-switching;
  }
}
xe-0/1/1 {
  unit 0 {
    family ethernet-switching;
  }
}
ge-0/1/2 {
  unit 0 {
    family ethernet-switching;
  }
}
ge-0/1/3 {
```

```
        unit 0 {
            family ethernet-switching;
        }
    }
}
    protocols {
        igmp-snooping {
            vlan all;
        }
        lldp {
            interface all;
        }
        lldp-med {
            interface all;
        }
        gvrp {
            interface xe-0/1/0.0;
        }
        rstp;
    }
    ethernet-switching-options {
        storm-control {
            interface all {
                level 50;
            }
        }
    }
    vlans {
        finance {
            vlan-id 100;
        }
        lab {
            vlan-id 200;
        }
        sales {
            vlan-id 300;
        }
    }
```

Configuring VLANS and GVRP on the Distribution Switch

CLI Quick Configuration

To quickly configure the **finance**, **lab**, and **sales** VLANs on the Distribution Switch and to enable GVRP on the trunk interface of the Distribution Switch, copy the following commands and paste them into the switch terminal window of the Distribution Switch:

```
[edit]
set vlans finance vlan-id 100
set vlans lab vlan-id 200
set vlans sales vlan-id 300
set interfaces xe-0/1/1 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/1/0 unit 0 family ethernet-switching port-mode trunk
set protocols gvrp interface xe-0/1/1.0
set protocols gvrp interface xe-0/1/0.0
```

Step-by-Step Procedure

To configure the three VLANs on the Distribution Switch, to configure the trunk interfaces, and to enable GVRP on the trunk interface of the Distribution Switch:

1. Configure the finance VLAN:

```
[edit]
user@Distribution-Switch# set vlans finance vlan–id 100
```

2. Configure the lab VLAN:

```
[edit]
user@Distribution-Switch# set vlans lab vlan–id 200
```

3. Configure the sales VLAN:

```
[edit]
user@Distribution-Switch# set vlans sales vlan–id 300
```

4. Configure the trunk interface to Access Switch A:

```
[edit]
user@Distribution-Switch# set interfaces xe-0/1/1 unit 0 familyethernet-switching
port-mode trunk
```

5. Configure the trunk interface to Access Switch B:

```
[edit]
user@Distribution-Switch# set interfaces xe-0/1/0 unit 0 family ethernet-switching
port-mode trunk
```

6. Enable GVRP on the trunk interface for **xe-0/1/1** :

```
[edit]
user@Distribution-Switch# set protocols gvrp interface xe-0/1/1.0
```

7. Enable GVRP on the trunk interface for **xe-0/1/0** :

```
[edit]
user@Distribution-Switch# set protocols gvrp interface xe-0/1/0.0
```

Results

Display the results of the configuration:

```
[edit]
  user@Distribution Switch-D #show
interfaces {
  ge-0/0/0 {
    unit 0 {
      family ethernet-switching;
    }
  }
  ge-0/0/1 {
    unit 0 {
      family ethernet-switching;
    }
  }
  ge-0/0/2 {
    unit 0 {
      family ethernet-switching;
    }
  }
```

```
    ge-0/0/3 {
        unit 0 {
            family ethernet-switching;
        }
    }
    xe-0/1/0 {
        unit 0 {
            family ethernet-switching {
                port-mode trunk;
            }
        }
    }
    ge-0/1/1 {
        unit 0 {
            family ethernet-switching;
        }
    }
    xe-0/1/1 {
        unit 0 {
            family ethernet-switching {
                port-mode trunk;
            }
        }
    }
    ge-0/1/2 {
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/1/3 {
        unit 0 {
            family ethernet-switching;
        }
    }
}
    protocols {
        igmp-snooping {
            vlan all;
        }
        lldp {
            interface all;
        }
        lldp-med {
            interface all;
        }
        gvrp {
            interface xe-0/1/0.0;
            interface xe-0/1/1.0;
        }
        rstp;
    }
    ethernet-switching-options {
        storm-control {
            interface all {
                level 50;
            }
```

```
    }
  }
  vlans {
    finance {
      vlan-id 100;
    }
    lab {
      vlan-id 300;
    }
    sales {
      vlan-id 300;
    }
  }
```

Verification

To confirm that the configuration is updating VLAN membership, perform these tasks:

- Verifying That GVRP Is Enabled on Access Switch A on page 1099
- Verifying That GVRP Is Updating VLAN Membership on Switch A on page 1099
- Verifying That GVRP Is Enabled on Access Switch B on page 1100
- Verifying That GVRP Is Updating VLAN Membership on Switch B on page 1100
- Verifying That GVRP Is Enabled on the Distribution Switch on page 1100
- Verifying That GVRP Is Updating VLAN Membership on the Distribution Switch on page 1101

Verifying That GVRP Is Enabled on Access Switch A

Purpose Verify that GVRP is enabled on the switch.

Action Show the GVRP configuration, using the **show gvrp** command:

```
user@Access-Switch-A> show gvrp

Global GVRP configuration
  GVRP status  : Enabled
  GVRP Timers (ms)
    Join       : 200
    Leave      : 600
    LeaveAll   : 10000
Interface Name      Protocol Status
--------------      ---------------
xe-0/1/1.0          Enabled
```

Meaning The results show that GVRP is enabled on the trunk interface of Switch A and that the default timers are used.

Verifying That GVRP Is Updating VLAN Membership on Switch A

Purpose To verify that GVRP is updating VLAN membership, display the Ethernet switching interfaces and associated VLANs that are active on switch A:

Action List Ethernet switching interfaces on the switch, using the **show ethernet-switching interfaces** command:

```
user@Access-Switch-A> show ethernet-switching interfaces
Interface   State    VLAN members           Blocking
ge-0/0/1.0  up       finance                unblocked
ge-0/0/2.0  up       lab                    unblocked
ge-0/0/3.0  up       sales                  unblocked
xe-0/1/1.0  up       finance                unblocked
                     lab                    unblocked
```

Meaning GVRP has automatically added **finance** and **lab** as VLAN members on the trunk interface, because they are being advertised by Access Switch B.

Verifying That GVRP Is Enabled on Access Switch B

Purpose Verify that GVRP is enabled on the switch.

Action Show the GVRP configuration:

```
user@Access-Switch-B> show gvrp

Global GVRP configuration
  GVRP status  : Enabled
  GVRP Timers (ms)
    Join       : 200
    Leave      : 600
    LeaveAll   : 10000
Interface Name      Protocol Status
--------------      ---------------
xe-0/1/0.0          Enabled
```

Meaning The results show that GVRP is enabled on the trunk interface of Switch B and that the default timers are used.

Verifying That GVRP Is Updating VLAN Membership on Switch B

Purpose To verify that GVRP is updating VLAN membership, display the Ethernet switching interfaces and associated VLANs that are active on switch B:

Action List Ethernet switching interfaces on the switch:

```
user@Access-Switch-B> show ethernet-switching interfaces
Interface   State    VLAN members           Blocking
ge-0/0/0.0  up       finance                unblocked
ge-0/0/1.0  up       lab                    unblocked
xe-0/1/1.0  up       finance                unblocked
                     lab                    unblocked
                     sales                  unblocked
```

Meaning GVRP has automatically added **finance**, **lab**, and **sales** as VLAN members on the trunk interface because they are being advertised by Access Switch A.

Verifying That GVRP Is Enabled on the Distribution Switch

Purpose Verify that GVRP is enabled on the switch.

Action Show the GVRP configuration:

```
user@Distribution-Switch> show gvrp

Global GVRP configuration
  GVRP status  : Enabled
  GVRP Timers (ms)
    Join       : 200
    Leave      : 600
    LeaveAll   : 10000
Interface Name      Protocol Status
--------------      ---------------
xe-0/1/0.0          Enabled
xe-0/1/1.0          Enabled
```

Verifying That GVRP Is Updating VLAN Membership on the Distribution Switch

Purpose To verify that GVRP is updating VLAN membership on the distribution switch, display the Ethernet switching interfaces and associated VLANs on the Distribution Switch:

Action List the Ethernet switching interfaces on the switch:

```
user@Distribution-Switch> show ethernet-switching interfaces
Interface   State    VLAN members           Blocking
xe-0/1/1.0  up       finance                unblocked
                     lab                    unblocked
                     sales                  unblocked
xe-0/1/0.0  up       finance                unblocked
                     lab                    unblocked
```

Meaning The Distribution Switch has two trunk interfaces. Interface **xe-0/1/1.0** connects the Distribution Switch to Access Switch A and is therefore updated to show that it is a member of all the VLANs that are active on Access Switch A. Any traffic for those VLANs will be passed on from the Distribution Switch to Access Switch A, through interface **xe-0/1/1.0**. Interface **xe-0/1/0.0** connects the Distribution Switch to Access Switch B and is updated to show that it is a member of the two VLANs that are active on Access Switch B. Thus, the Distribution Switch sends traffic for **finance** and **lab** to both Access Switch A and Access Switch B. But the Distribution Switch sends traffic for **sales** only to Access Switch A.

Related Documentation

- Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch on page 1063
- Understanding Bridging and VLANs on J-EX Series Switches on page 1041

Example: Configuring Redundant Trunk Links for Faster Recovery

Simplify the convergence configuration in a typical enterprise network by configuring a primary link and a secondary link on trunk ports. If the primary link fails, the secondary link automatically takes over without waiting for normal STP convergence.

This example describes how to create a redundant trunk group:

- Requirements on page 1102
- Overview and Topology on page 1102

Requirements

This example uses the following hardware and software components:

- Two J-EX4200 distribution switches
- One J-EX4200 access switch

Before you configure the redundant trunk links network on the access and distribution switches, be sure you have:

- Installed the access switch. See Installing and Connecting a J-EX4200 Switch.
- Installed the two distribution switches. See Installing and Connecting a J-EX4200 Switch.
- Performed the initial switch configuration. See “Connecting and Configuring a J-EX Series Switch (J-Web Procedure)” on page 163.

Overview and Topology

This example shows a simple configuration to illustrate the basic steps for creating a redundant trunk group.

Configuring redundant trunk links places the primary link and the secondary link in a redundant group. However, a primary link need not be configured. If a primary link is not specified, the software compares the two links and selects the link with the highest port number as the active link. For example, if the two interfaces are **ge-0/1/0** and **ge-0/1/1**, the software assigns **ge-0/1/1** as the active link..

Whether a primary link is specified as the active link, or whether it is calculated by the software, traffic is handled in the same manner. Traffic passes through the active link but is blocked on the secondary link. If the active link goes down or is disabled administratively, the secondary link becomes active and begins forwarding traffic. However, there is a difference between the behavior of a primary, active link and an active link that is calculated to be active by the software. If an active link goes down, the secondary link begins forwarding traffic. If the old, active link comes up again, the following occurs:

- If the old, active link was configured as the primary link, then it resumes the role of active link and the other link is blocked. An interface configured as primary continues to carry with it the primary role whenever it becomes active.
- If no primary link was configured, and the active link was calculated by the software when the redundant group was formed, then the old, active link will not preempt the other interface (new active).

NOTE: The Junos OS for J-EX Series switches does not allow an interface to be in a redundant trunk group and in an STP topology at the same time.

Figure 30 on page 1103 displays an example topology containing three switches. Switch 1 and Switch 2 make up the distribution layer, and Switch 3 makes up the access layer. Switch 3 is connected to the distribution layer through trunk ports **ge-0/0/9.0** (Link 1) and **ge-0/0/10.0** (Link 2).

Table 146 on page 1103 lists the components used in this redundant trunk group.

Figure 30: Topology for Configuring the Redundant Trunk Links

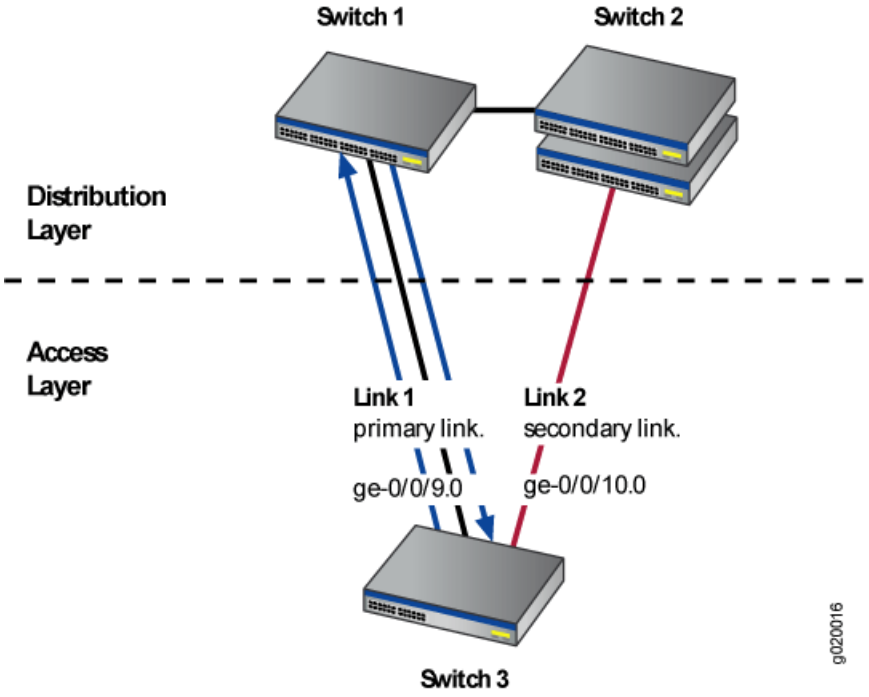

Table 146: Components of the Redundant Trunk Link Topology

Property	Settings
Switch hardware	• Switch 1–1 J-EX4200 distribution switch • Switch 2–1 J-EX4200 distribution switch • Switch 3–1 J-EX4200 access switch
Trunk port interfaces	On Switch 3 (access switch): **ge-0/0/9.0** and **ge-0/0/10.0**
Redundant trunk group	**group1**

This configuration example creates a redundant trunk group called **group1** on Switch 3. The trunk ports **ge-0/0/9.0** and **ge-0/0/10.0** are the two links in **group1**. The trunk port **ge-0/0/9.0** will be configured administratively as the primary link. The trunk port **ge-0/0/10.0** will be the secondary link.

Configuration

CLI Quick Configuration

To quickly configure the redundant trunk group **group1** on Switch 3, copy the following commands and paste them into the switch terminal window:

```
[edit]
set ethernet-switching-options redundant-trunk-group group-name group1
set ethernet-switching-options redundant-trunk-group group-name group1 interface ge-0/0/9.0
primary
set ethernet-switching-options redundant-trunk-group group-name group1 interface ge-0/0/10.0
```

Step-by-Step Procedure

Configure the redundant trunk group **group1** on Switch 3 and specify the primary and secondary links.

1. Configure the redundant trunk group **group1**:

   ```
   [edit ethernet-switching-options]
   user@switch# set redundant-trunk-group group-name group1
   ```

2. Configure the trunk port **ge-0/0/9.0** as the primary link and **ge-0/0/10** as the secondary link:

   ```
   [edit ethernet-switching-options]
   user@switch# set redundant-trunk-group group-name group1 interface ge-0/0/9.0
   primary
   user@switch# set redundant-trunk-group group-name group1 interface ge-0/0/10.0
   ```

Results

Display the results of the configuration:

```
user@switch# show
  ethernet-switching-options {
    redundant-trunk-group {
      group-name group1 {
        interface ge-0/0/9.0 primary;
        interface ge-0/0/10.0;
      }
    }
  }
}
```

Verification

Verify that the redundant trunk group **group1** has been created and is operating properly:

- Verifying That the Redundant Group Has Been Created on page 1104

Verifying That the Redundant Group Has Been Created

Purpose

Verify that the redundant trunk group **group1** has been created on the switch and that trunk ports are members of the redundant trunk group.

Action

List all redundant trunk groups configured on the switch:

```
user@switch> show redundant-trunk-group group1
Redundant-trunk-group: group1
Interfaces           : ge-0/0/9.0  (P) , DOWN
                     : ge-0/0/10.0 (A) , UP
Bandwidth            : 1000 Mbps, 1000 Mbps
Last Time of Flap    : 1970-01-01 00:19:12 UTC (00:00:06 ago), Never
#Flaps               : 1, 0
```

Meaning The **show redundant-trunk-group** command lists all redundant trunk groups configured on the switch and which trunk links are members of the group. For this configuration example, the output shows that the redundant trunk group **group1** is configured on the switch. The **(P)** beside trunk port **ge-0/0/9.0** indicates that it is configured as the primary link. The **(A)** beside the **ge-0/0/10.0** trunk port indicates that it is the active link.

Related Documentation

- Understanding Redundant Trunk Links on J-EX Series Switches on page 1049

Example: Setting Up Q-in-Q Tunneling on J-EX Series Switches

Service providers can use Q-in-Q tunneling to transparently pass Layer 2 VLAN traffic from a customer site, through the service provider network, to another customer site without removing or changing the customer VLAN tags or class-of-service (CoS) settings. You can configure Q-in-Q tunneling on J-EX Series switches.

This example describes how to set up Q-in-Q:

Requirements

This example requires one J-EX Series switch.

Before you begin setting up Q-in-Q tunneling, make sure you have created and configured the necessary customer VLANs. See “Configuring VLANs for J-EX Series Switches (CLI Procedure)” on page 1136 or “Configuring VLANs for J-EX Series Switches (J-Web Procedure)” on page 1133.

Overview and Topology

In this service provider network, there are multiple customer VLANs mapped to one service VLAN.

Table 147 on page 1105 lists the settings for the example topology.

Table 147: Components of the Topology for Setting Up Q-in-Q Tunneling

Interface	Description
ge-0/0/11.0	Tagged S-VLAN trunk port
ge-0/0/12.0	Untagged customer-facing access port
ge-0/0/13.0	Untagged customer-facing access port
ge-0/0/14.0	Tagged S-VLAN trunk port

Configuration

CLI Quick Configuration

To quickly create and configure Q-in-Q tunneling, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans qinqvlan vlan-id 4001
set vlans qinqvlan dot1q-tunneling customer-vlans 1-100
set vlans qinqvlan dot1q-tunneling customer-vlans 201-300
set interfaces ge-0/0/11 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/11 unit 0 family ethernet-switching vlan members 4001
set interfaces ge-0/0/12 unit 0 family ethernet-switching port-mode access
set interfaces ge-0/0/12 unit 0 family ethernet-switching vlan members 4001
set interfaces ge-0/0/13 unit 0 family ethernet-switching port-mode access
set interfaces ge-0/0/13 unit 0 family ethernet-switching vlan members 4001
set interfaces ge-0/0/14 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/14 unit 0 family ethernet-switching vlan members 4001
set ethernet-switching-options dot1q-tunneling ether-type 0x9100
```

Step-by-Step Procedure

To configure Q-in-Q tunneling:

1. Set the VLAN ID for the S-VLAN:

```
[edit vlans]
user@switch# set qinqvlan vlan-id 4001
```

2. Enable Q-in-Q tuennling and specify the customer VLAN ranges:

```
[edit vlans]
user@switch# set qinqvlan dot1q-tunneling customer-vlans 1-100
user@switch# set qinqvlan dot1q-tunneling customer-vlans 201-300
```

3. Set the port mode and VLAN information for the interfaces:

```
[edit interfaces]
user@switch# set ge-0/0/11 unit 0 family ethernet-switching port-mode trunk
user@switch# set ge-0/0/11 unit 0 family ethernet-switching vlan members 4001
user@switch# set ge-0/0/12 unit 0 family ethernet-switching port-mode access
user@switch# set ge-0/0/12 unit 0 family ethernet-switching vlan members 4001
user@switch# set ge-0/0/13 unit 0 family ethernet-switching port-mode access
user@switch# set ge-0/0/13 unit 0 family ethernet-switching vlan members 4001
user@switch# set ge-0/0/14 unit 0 family ethernet-switching port-mode trunk
user@switch# set ge-0/0/14 unit 0 family ethernet-switching vlan members 4001
```

4. Set the Q-in-Q Ethertype value:

```
[edit]
user@switch# set ethernet-switching-options dot1q-tunneling ether-type 0x9100
```

Results

Check the results of the configuration:

```
user@switch> show configuration vlans qinqvlan
vlan-id 4001;
   dot1q-tunneling {
   customer-vlans [ 1-100 201-300 ];
}
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying That Q-in-Q Tunneling Was Enabled on page 1107

Verifying That Q-in-Q Tunneling Was Enabled

Purpose Verify that Q-in-Q tunneling was properly enabled on the switch.

Action Use the **show vlans** command:

```
user@switch> show vlans qinqvlan extensive
VLAN: qinqvlan, Created at: Thu Sep 18 07:17:53 2008
802.1Q Tag: 4001, Internal index: 18, Admin State: Enabled, Origin: Static
Dot1q Tunneling Status: Enabled
Customer VLAN ranges:
                     1-100
                     201-300
Protocol: Port Mode
Number of interfaces: Tagged 2 (Active = 0), Untagged  4 (Active = 0)
      ge-0/0/11.0, tagged, trunk
      ge-0/0/14.0, tagged, trunk
      ge-0/0/12.0, untagged, access
      ge-0/0/13.0, untagged, access
```

Meaning The output indicates that Q-in-Q tunneling is enabled and that the VLAN is tagged and shows the associated customer VLANs.

Related Documentation

- Configuring Q-in-Q Tunneling (CLI Procedure) on page 1144

Example: Configuring a Private VLAN on a J-EX Series Switch

For security reasons, it is often useful to restrict the flow of broadcast and unknown unicast traffic and to even limit the communication between known hosts. The private VLAN (PVLAN) feature on J-EX Series switches allow an administrator to split a broadcast domain into multiple isolated broadcast subdomains, essentially putting a VLAN inside a VLAN.

This example describes how to create a private VLAN primary VLAN and secondary VLANs:

NOTE: Configuring a voice over IP (VoIP) VLAN on PVLAN interfaces is not supported.

- Requirements on page 1108
- Overview and Topology on page 1108
- Configuration on page 1108
- Verification on page 1110

Requirements

This example requires one J-EX Series switch.

Before you begin configuring a private VLAN, make sure you have created and configured the necessary VLAN. See “Configuring VLANs for J-EX Series Switches (CLI Procedure)” on page 1136 or “Configuring VLANs for J-EX Series Switches (J-Web Procedure)” on page 1133.

Overview and Topology

In a large office with multiple buildings and VLANs, you might need to isolate some workgroups or other endpoints for security reasons or to partition the broadcast domain. This configuration example shows a simple topology to illustrate how to create a private VLAN with one primary VLAN and two community VLANs, one for HR and one for finance, as well as two isolated ports for the mail server and the backup server.

Table 148 on page 1108 lists the settings for the example topology.

Table 148: Components of the Topology for Configuring a Private VLAN

Interface	Description
ge-0/0/0.0	Primary VLAN (**pvlan**) trunk interface
ge-0/0/11.0	User 1, HR Community (**hr-comm**)
ge-0/0/12.0	User 2, HR Community (**hr-comm**)
ge-0/0/13.0	User 3, Finance Community (**finance-comm**)
ge-0/0/14.0	User 4, Finance Community (**finance-comm**)
ge-0/0/15.0	Mail server, Isolated (**isolated**)
ge-0/0/16.0	Backup server, Isolated (**isolated**)
ge-1/0/0.0	Primary VLAN (**pvlan**) trunk interface

Configuration

CLI Quick Configuration

To quickly create and configure a private VLAN, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans pvlan vlan-id 1000
set interfaces ge-0/0/0 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/0 unit 0 family ethernet-switching vlan members pvlan
set interfaces ge-1/0/0 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-1/0/0 unit 0 family ethernet-switching vlan members pvlan
set interfaces ge-0/0/11 unit 0 family ethernet-switching port-mode access
set interfaces ge-0/0/12 unit 0 family ethernet-switching port-mode access
set interfaces ge-0/0/13 unit 0 family ethernet-switching port-mode access
set interfaces ge-0/0/14 unit 0 family ethernet-switching port-mode access
```

```
set interfaces ge-0/0/15 unit 0 family ethernet-switching port-mode access
set interfaces ge-0/0/16 unit 0 family ethernet-switching port-mode access
set vlans pvlan no-local-switching
set vlans pvlan interface ge-0/0/0.0
set vlans pvlan interface ge-1/0/0.0
set vlans hr-comm interface ge-0/0/11.0
set vlans hr-comm interface ge-0/0/12.0
set vlans finance-comm interface ge-0/0/13.0
set vlans finance-comm interface ge-0/0/14.0
set vlans hr-comm primary-vlan pvlan
set vlans finance-comm primary-vlan pvlan
```

Step-by-Step Procedure

To configure the private VLAN:

1. Set the VLAN ID for the primary VLAN:

```
[edit vlans]
user@switch# set pvlan vlan-id 1000
```

2. Set the interfaces and port modes:

```
[edit interfaces]
user@switch# set ge-0/0/0 unit 0 family ethernet-switching port-mode trunk

user@switch# set ge-0/0/0 unit 0 family ethernet-switching vlan members pvlan

user@switch# set ge-1/0/0 unit 0 family ethernet-switching port-mode trunk

user@switch# set ge-1/0/0 unit 0 family ethernet-switching vlan members pvlan

user@switch# set ge-0/0/11 unit 0 family ethernet-switching port-mode access

user@switch# set ge-0/0/12 unit 0 family ethernet-switching port-mode access

user@switch# set ge-0/0/13 unit 0 family ethernet-switching port-mode access

user@switch# set ge-0/0/14 unit 0 family ethernet-switching port-mode access

user@switch# set ge-0/0/15 unit 0 family ethernet-switching port-mode access

user@switch# set ge-0/0/16 unit 0 family ethernet-switching port-mode access
```

3. Set the primary VLAN to have no local switching:

NOTE: The primary VLAN must be a tagged VLAN.

```
[edit vlans]
user@switch# set pvlan no-local-switching
```

4. Add the trunk interfaces to the primary VLAN:

```
[edit vlans]
user@switch# set pvlan interface ge-0/0/0.0

user@switch# set pvlan interface ge-1/0/0.0
```

5. For each secondary VLAN, configure access interfaces:

NOTE: The secondary VLANs must be untagged VLANs.

```
[edit vlans]
user@switch# set hr-comm interface ge-0/0/11.0

user@switch# set hr-comm interface ge-0/0/12.0

user@switch# set finance-comm interface ge-0/0/13.0

user@switch# set finance-comm interface ge-0/0/14.0
```

6. For each community VLAN, set the primary VLAN:

```
[edit vlans]
user@switch# set hr-comm primary-vlan pvlan

user@switch# set finance-comm primary-vlan pvlan
```

7. Add each isolated interface to the primary VLAN:

```
[edit vlans]
user@switch# set pvlan interface ge-0/0/15.0

user@switch# set pvlan interface ge-0/0/16.0
```

Results Check the results of the configuration:

```
user@switch> show configuration vlans
finance-comm {
    interface {
        ge-0/0/13.0;
        ge-0/0/14.0;
    }
    primary-vlan pvlan;
}
    hr-comm {
        interface {
            ge-0/0/11.0;
            ge-0/0/12.0;
        }
        primary-vlan pvlan;
    }
    pvlan {
        vlan-id 1000;
        interface {
            ge-0/0/15.0;
            ge-0/0/16.0;
            ge-0/0/0.0;
            ge-1/0/0.0;
        }
        no-local-switching;
    }
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying the Private VLAN and Secondary VLANs Were Created on page 1111

Verifying the Private VLAN and Secondary VLANs Were Created

Purpose Verify that the primary VLAN and secondary VLANs were properly created on the switch.

Action Use the **show vlans** command:

```
user@switch> show vlans pvlan extensive
VLAN: pvlan, Created at: Tue Sep 16 17:59:47 2008
802.1Q Tag: 1000, Internal index: 18, Admin State: Enabled, Origin: Static
Private VLAN Mode: Primary
Protocol: Port Mode
Number of interfaces: Tagged 2 (Active = 0), Untagged  6 (Active = 0)
      ge-0/0/0.0, tagged, trunk
      ge-0/0/11.0, untagged, access
      ge-0/0/12.0, untagged, access
      ge-0/0/13.0, untagged, access
      ge-0/0/14.0, untagged, access
      ge-0/0/15.0, untagged, access
      ge-0/0/16.0, untagged, access
      ge-1/0/0.0, tagged, trunk
Secondary VLANs: Isolated 2, Community  2
  Isolated VLANs :
      __pvlan_pvlan_ge-0/0/15.0__
      __pvlan_pvlan_ge-0/0/16.0__
  Community VLANs :
      finance-comm
      hr-comm

user@switch> show vlans hr-comm extensive
VLAN: hr-comm, Created at: Tue Sep 16 17:59:47 2008
Internal index: 22, Admin State: Enabled, Origin: Static
Private VLAN Mode: Community, Primary VLAN: pvlan
Protocol: Port Mode
Number of interfaces: Tagged 2 (Active = 0), Untagged  2 (Active = 0)
      ge-0/0/0.0, tagged, trunk
      ge-0/0/11.0, untagged, access
      ge-0/0/12.0, untagged, access
      ge-1/0/0.0, tagged, trunk

user@switch> show vlans finance-comm extensive
VLAN: finance-comm, Created at: Tue Sep 16 17:59:47 2008
Internal index: 21, Admin State: Enabled, Origin: Static
Private VLAN Mode: Community, Primary VLAN: pvlan
Protocol: Port Mode
Number of interfaces: Tagged 2 (Active = 0), Untagged  2 (Active = 0)
      ge-0/0/0.0, tagged, trunk
      ge-0/0/13.0, untagged, access
      ge-0/0/14.0, untagged, access
      ge-1/0/0.0, tagged, trunk

user@switch> show vlans __pvlan_pvlan_ge-0/0/15.0__ extensive
VLAN: __pvlan_pvlan_ge-0/0/15.0__, Created at: Tue Sep 16 17:59:47 2008
Internal index: 19, Admin State: Enabled, Origin: Static
Private VLAN Mode: Isolated, Primary VLAN: pvlan
Protocol: Port Mode
Number of interfaces: Tagged 2 (Active = 0), Untagged  1 (Active = 0)
      ge-0/0/0.0, tagged, trunk
      ge-0/0/15.0, untagged, access
      ge-1/0/0.0, tagged, trunk

user@switch> show vlans __pvlan_pvlan_ge-0/0/16.0__ extensive
```

```
VLAN: __pvlan_pvlan_ge-0/0/16.0__, Created at: Tue Sep 16 17:59:47 2008
Internal index: 20, Admin State: Enabled, Origin: Static
Private VLAN Mode: Isolated, Primary VLAN: pvlan
Protocol: Port Mode
Number of interfaces: Tagged 2 (Active = 0), Untagged  1 (Active = 0)
      ge-0/0/0.0, tagged, trunk
      ge-0/0/16.0, untagged, access
      ge-1/0/0.0, tagged, trunk
```

Meaning The output shows that the primary VLAN was created and identifies the interfaces and secondary VLANs associated with it.

Related Documentation

- Creating a Private VLAN (CLI Procedure) on page 1143

Example: Using Virtual Routing Instances to Route Among VLANs on J-EX Series Switches

Virtual routing instances allow each J-EX Series switch to have multiple routing tables on a device. With virtual routing instances, you can segment your network to isolate traffic without setting up additional devices.

This example describes how to create virtual routing instances:

Requirements

This example uses the following hardware and software components:

- One J-EX Series switch

Before you create the virtual routing instances, make sure you have:

- Configured the necessary VLANs. See “Configuring VLANs for J-EX Series Switches (CLI Procedure)” on page 1136 or “Configuring VLANs for J-EX Series Switches (J-Web Procedure)” on page 1133.

Overview and Topology

In a large office, you may need multiple VLANs to properly manage your traffic. This configuration example shows a simple topology to illustrate how to connect a single J-EX Series switch with a virtual routing instance for each of two VLANs, enabling traffic to pass between those VLANs.

In the example topology, the LAN is segmented into two VLANs, each associated with an interface and a routing instance on the J-EX Series switch.

Configuration

CLI Quick Configuration

To quickly create and configure virtual routing instances, copy the following commands and paste them into the switch terminal window:

```
[edit]
set interfaces ge-0/0/3 vlan-tagging
set interfaces ge-0/0/3 unit 0 vlan-id 1030 family inet address 103.1.1.1/24
set interfaces ge-0/0/3 unit 1 vlan-id 1031 family inet address 103.1.1.1/24
set routing-instances r1 instance-type virtual-router
set routing-instances r1 interface ge-0/0/1.0
set routing-instances r1 interface ge-0/0/3.0
set routing-instances r2 instance-type virtual-router
set routing-instances r2 interface ge-0/0/2.0
set routing-instances r2 interface ge-0/0/3.1
```

Step-by-Step Procedure

To configure virtual routing instances:

1. Create a VLAN-tagged interface:

   ```
   [edit]
   user@switch# set interfaces ge-0/0/3 vlan-tagging
   ```

2. Create two subinterfaces, on the interface, one for each routing instance:

   ```
   [edit]
   user@switch# set interfaces ge-0/0/3 unit 0 vlan-id 1030 family inet address 103.1.1.1/24

   user@switch# set interfaces ge-0/0/3 unit 1 vlan-id 1031 family inet address 103.1.1.1/24
   ```

3. Create two virtual routers:

   ```
   [edit]
   user@switch# set routing-instances r1 instance-type virtual-router
   user@switch# set routing-instances r2 instance-type virtual-router
   ```

4. Set the interfaces for the virtual routers:

   ```
   [edit]
   user@switch# set routing-instances r1 interface ge-0/0/1.0

   user@switch# set routing-instances r1 interface ge-0/0/3.0

   user@switch# set routing-instances r2 interface ge-0/0/2.0

   user@switch# set routing-instances r2 interface ge-0/0/3.1
   ```

Results

Check the results of the configuration:

```
user@switch> show configuration
interfaces {
    ge-0/0/1 {
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/0/2 {
        unit 0 {
            family ethernet-switching;
        }
```

```
}
ge-0/0/3 {
  vlan-tagging;
  unit 0 {
    vlan-id 1030;
    family inet {
      address 103.1.1.1/24;
    }
  }
  unit 1 {
    vlan-id 1031;
    family inet {
      address 103.1.1.1/24;
    }
  }
}
routing-instances {
  r1 {
    instance-type virtual-router;
    interface ge-0/0/1.0;
    interface ge-0/0/3.0;
  }
  r2 {
    instance-type virtual-router;
    interface ge-0/0/2.0;
    interface ge-0/0/3.1;
  }
}
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying That the Routing Instances Were Created on page 1114

Verifying That the Routing Instances Were Created

Purpose Verify that the virtual routing instances were properly created on the switch.

Action Use the **show route instance** command:

```
user@switch> show route instance
Instance             Type
        Primary RIB                                    Active/holddown/hidden
master               forwarding
        inet.0                                         3/0/0

r1                   virtual-router
        r1.inet.0                                      1/0/0

r2                   virtual-router
        r2.inet.0                                      1/0/0
```

Meaning Each routing instance created is displayed, along with its type, information about whether it is active or not, and its primary routing table.

Related Documentation

- Configuring Virtual Routing Instances (CLI Procedure) on page 1142

Example: Configuring Automatic VLAN Administration Using MVRP on J-EX Series Switches

As a network expands and the number of clients and VLANs increases, VLAN administration becomes complex and the task of efficiently configuring VLANs on multiple J-EX Series switches becomes increasingly difficult. To automate VLAN administration, you can enable Multiple VLAN Registration Protocol (MVRP) on the network.

MVRP can also be used to dynamically create VLANs, further simplifying the network overhead required to statically configure VLANs.

NOTE: Only trunk interfaces can be enabled for MVRP.

This example describes how to use MVRP to automate administration of VLAN membership changes within your network and how to use MVRP to dynamically create VLANs:

Requirements

This example uses the following hardware and software components:

- Two J-EX Series access switches
- One J-EX Series distribution switch

Overview and Topology

MVRP is used to manage dynamic VLAN registration in a LAN. It can also be used to dynamically create VLANs.

This example uses MVRP to dynamically create VLANs on the switching network. You can disable dynamic VLAN creation and create VLANs statically, if desired. Enabling MVRP on the trunk interface of each switch in your switching network ensures that the active VLAN information for the switches in the network is propagated to each switch through the trunk interfaces, assuming dynamic VLAN creation is enabled for MVRP.

MVRP ensures that the VLAN membership information on the trunk interface is updated as the switch's access interfaces become active or inactive in the configured VLANs in a static or dynamic VLAN creation setup.

You do not need to explicitly bind a VLAN to the trunk interface. When MVRP is enabled, the trunk interface advertises all the VLANs that are active (bound to access interfaces) on that switch. An MVRP-enabled trunk interface does not advertise VLANs that have been configured on the switch but that are not currently bound to an access interface. Thus, MVRP provides the benefit of reducing network overhead—by limiting the scope of broadcast, unknown unicast, and multicast (BUM) traffic to interested devices only.

When VLAN access interfaces become active or inactive, MVRP ensures that the updated information is advertised on the trunk interface. Thus, in this example, distribution Switch C does not forward traffic to inactive VLANs.

This example shows a network with three VLANs: **finance**, **sales**, and **lab**.

Access Switch A has been configured to support all three VLANS and all three VLANS are active, bound to interfaces that are connected to personal computers:

- **ge-0/0/1**—Connects PC1 as a member of **finance**, VLAN ID 100
- **ge-0/0/2**—Connects PC2 as a member of **lab**, VLAN ID 200
- **ge-0/0/3**—Connects PC3 as a member of **sales**, VLAN ID 300

Access Switch B has also been configured to support three VLANS. However, currently only two VLANs are active, bound to interfaces that are connected to personal computers:

- **ge-0/0/0**—Connects PC4 as a member of **finance**, VLAN ID 100
- **ge-0/0/1**—Connects PC5 as a member of **lab**, VLAN ID 200

Distribution Switch C learns the VLANs dynamically using MVRP through the connection to the access switches. Distribution Switch C has two trunk interfaces:

- **xe-0/1/1**—Connects the switch to access Switch A.
- **xe-0/1/0**—Connects the switch to access Switch B.

Figure 31 on page 1117 shows MVRP configured on two access switches and one distribution switch.

Figure 31: MVRP Configured on Two Access Switches and One Distribution Switch for Automatic VLAN Administration

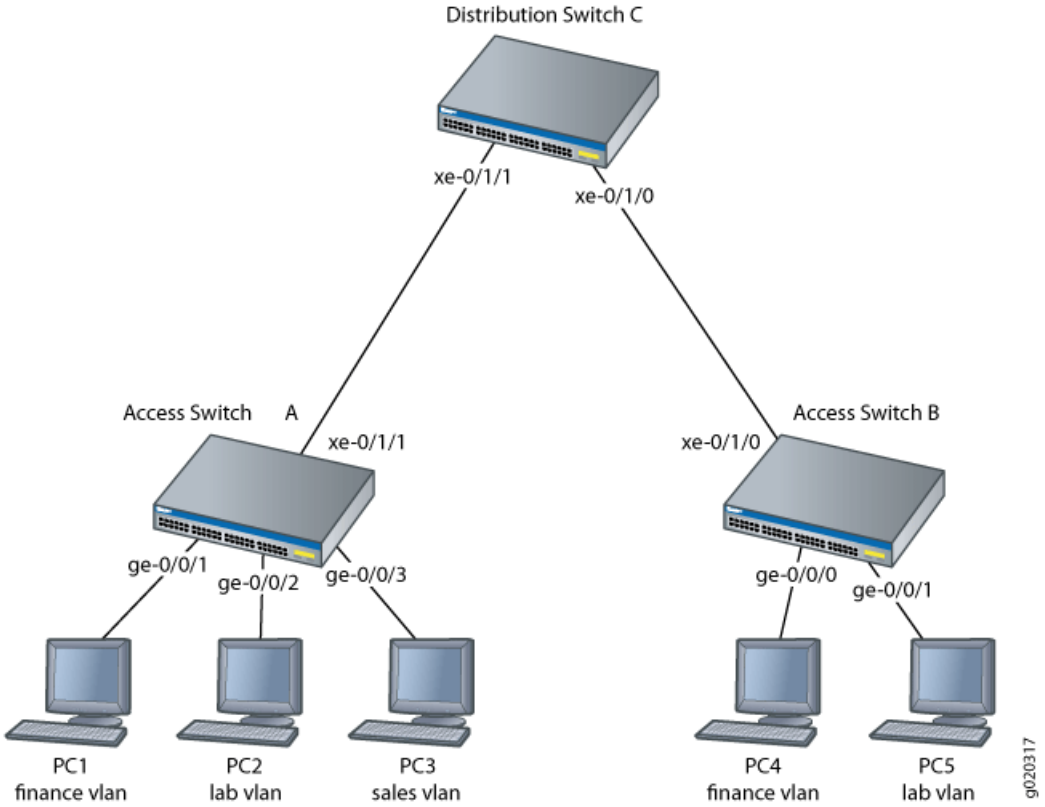

Table 149 on page 1117 explains the components of the example topology.

Table 149: Components of the Network Topology

Property	Settings
Switch hardware	• Access Switch A • Access Switch B • Distribution Switch C
VLAN names and tag IDs	**finance**, tag **100** **lab**, tag **200** **sales**, tag **300**

Table 149: Components of the Network Topology *(continued)*

Property	Settings
Interfaces	Access Switch A interfaces: • **ge-0/0/1**—Connects PC1 to access Switch A. • **ge-0/0/2**—Connects PC2 to access Switch A. • **ge-0/0/3**—Connects PC3 to access Switch A. • **xe-0/1/1**—Connects access Switch A to distribution Switch C (trunk). Access Switch B interfaces: • **ge-0/0/0**—Connects PC4 to access Switch B. • **ge-0/0/1**—Connects PC5 to access Switch B. • **xe-0/1/0**—Connects access Switch B to distribution Switch C. (trunk) Distribution Switch C interfaces: • **xe-0/1/1**—Connects distribution Switch C to access Switch A. (trunk) • **xe-0/1/0**—Connects distribution Switch C to access Switch B. (trunk)

Configuring VLANs and MVRP on Access Switch A

To configure VLANs on the switch, bind access interfaces to the VLANs, and enable MVRP on the trunk interface of access Switch A, perform these tasks:

CLI Quick Configuration

To quickly configure access Switch A for MVRP, copy the following commands and paste them into the switch terminal window of Switch A:

```
[edit]
set vlans finance vlan-id 100
set vlans lab vlan-id 200
set vlans sales vlan-id 300
set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan members finance
set interfaces ge-0/0/2 unit 0 family ethernet-switching vlan members lab
set interfaces ge-0/0/3 unit 0 family ethernet-switching vlan members sales
set interfaces xe-0/1/1 unit 0 family ethernet-switching port-mode trunk
set protocols mvrp interface xe-0/1/1.0
```

NOTE: As we recommend as a best practice, default MVRP timers are used in this example. The default values associated with each MVRP timer are: 200 ms for the join timer, 1000 ms for the leave timer, and 10000 ms for the leaveall timer. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.

Step-by-Step Procedure

To configure access Switch A for MVRP:

1. Configure the finance VLAN:

```
[edit]
user@Access-Switch-A# set vlans finance vlan-id 100
```

2. Configure the lab VLAN:

```
[edit]
user@Access-Switch-A# set vlans lab vlan–id 200
```

3. Configure the sales VLAN:

```
[edit]
user@Access-Switch-A# set vlans sales vlan–id 300
```

4. Configure an Ethernet interface as a member of the finance VLAN:

```
[edit]
user@Access-Switch-A# set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan
members finance
```

5. Configure an Ethernet interface as a member of the lab VLAN:

```
[edit]
user@Access-Switch-A# set interfaces ge-0/0/2 unit 0 family ethernet-switching vlan
members lab
```

6. Configure an Ethernet interface as a member of the sales VLAN:

```
[edit]
user@Access-Switch-A# set interfaces ge-0/0/3 unit 0 family ethernet-switching vlan
members sales
```

7. Configure a trunk interface:

```
[edit]
user@Access-Switch-A# set interfaces xe-0/1/1 unit 0 family ethernet-switching
port-mode trunk
```

8. Enable MVRP on the trunk interface:

```
[edit]
user@Access-Switch-A# set protocols mvrp interface xe-0/1/1.0
```

Results

Check the results of the configuration:

```
[edit]
  user@switch# show
  interfaces {
  ge-0/0/1 {
    unit 0 {
      family ethernet-switching {
        vlan {
          members finance;
        }
      }
    }
  }
  ge-0/0/2 {
```

```
        unit 0 {
          family ethernet-switching {
            vlan {
              members lab;
            }
          }
        }
      }
      ge-0/0/3 {
        unit 0 {
          family ethernet-switching {
            vlan {
              members sales;
            }
          }
        }
      }
      xe-0/1/1 {
        unit 0 {
          family ethernet-switching {
            port-mode trunk;
          }
        }
      }
    }

    protocols {
      mvrp {
        interface xe-0/1/1.0;
      }
    }

    vlans {
      finance {
        vlan-id 100;
      }
      lab {
        vlan-id 200;
      }
      sales {
        vlan-id 300;
        }
      }
```

Configuring VLANs and MVRP on Access Switch B

To configure three VLANs on the switch, bind access interfaces for PC4 and PC5 to the VLANs, and enable MVRP on the trunk interface of access Switch B, perform these tasks:

CLI Quick Configuration

To quickly configure Access Switch B for MVRP, copy the following commands and paste them into the switch terminal window of Switch B:

```
[edit]
set vlans finance vlan-id 100
set vlans lab vlan-id 200
set vlans sales vlan-id 300
set interfaces ge-0/0/0 unit 0 family ethernet-switching vlan members finance
```

```
set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan members lab
set interfaces xe-0/1/0 unit 0 family ethernet-switching port-mode trunk
set protocols mvrp interface xe-0/1/0.0
```

Step-by-Step Procedure

To configure access Switch B for MVRP:

1. Configure the finance VLAN:

```
[edit]
user@Access-Switch-B# set vlans finance vlan-id 100
```

2. Configure the lab VLAN:

```
[edit]
user@Access-Switch-B# set vlans lab vlan–id 200
```

3. Configure the sales VLAN:

```
[edit]
user@Access-Switch-B# set vlans sales vlan–id 300
```

4. Configure an Ethernet interface as a member of the finance VLAN:

```
[edit]
user@Access-Switch-B# set interfaces ge-0/0/0 unit 0 family ethernet-switching vlan
members finance
```

5. Configure an Ethernet interface as a member of the lab VLAN:

```
[edit]
user@Access-Switch-B# set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan
members lab
```

6. Configure a trunk interface:

```
user@Access-Switch-B# set interfaces xe-0/1/0 unit 0 family ethernet-switching
port-mode trunk
```

7. Enable MVRP on the trunk interface:

```
[edit]
user@Access-Switch-B# set protocols mvrp xe-0/1/0.0
```

NOTE: As we recommend as a best practice, default MVRP timers are used in this example. The default values associated with each MVRP timer are: 200 ms for the join timer, 1000 ms for the leave timer, and 10000 ms for the leaveall timer. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.

Results

Check the results of the configuration:

```
[edit]
user@Access-Switch-B# show
interfaces {
  ge-0/0/0 {
    unit 0 {
      family ethernet-switching {
        vlan {
```

```
            members finance;
          }
        }
      }
    }
    ge-0/0/1 {
      unit 0 {
        family ethernet-switching {
          vlan {
            members lab;
          }
        }
      }
    }
    xe-0/1/0 {
      unit 0 {
        family ethernet-switching {
          port-mode trunk;
        }
      }
    }
    }
}

protocols {
  mvrp {
    interface xe-0/1/0.0;
  }
}
vlans {
  finance {
    vlan-id 100;
  }
  lab {
    vlan-id 200;
  }
  sales {
    vlan-id 300;
  }
}
```

Configuring VLANS and MVRP on Distribution Switch C

CLI Quick Configuration

To quickly configure distribution Switch C for MVRP, copy the following commands and paste them into the switch terminal window of distribution Switch C:

```
[edit]
set interfaces xe-0/1/1 unit 0 family ethernet-switching port-mode trunk
set interfaces xe-0/1/0 unit 0 family ethernet-switching port-mode trunk
set protocols mvrp interface xe-0/1/1.0
set protocols mvrp interface xe-0/1/0.0
```

Step-by-Step Procedure

To configure distribution Switch C for MVRP:

1. Configure the trunk interface to access Switch A:

```
[edit]
user@Distribution-Switch-C# set interfaces xe-0/1/1 unit 0 family ethernet-switching
port-mode trunk
```

2. Configure the trunk interface to access Switch B:

```
[edit]
user@Distribution-Switch-C# set interfaces xe-0/1/0 unit 0 family ethernet-switching
port-mode trunk
```

3. Enable MVRP on the trunk interface for **xe-0/1/1** :

```
[edit]
user@Distribution-Switch-C# set protocols mvrp interface xe-0/1/1.0
```

4. Enable MVRP on the trunk interface for **xe-0/1/0** :

```
[edit]
user@Distribution-Switch-C# set protocols mvrp interface xe-0/1/0.0
```

Results

Check the results of the configuration:

```
[edit]
user@Distribution Switch-D# show
interfaces {
  xe-0/1/0 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
      }
    }
  }
  xe-0/1/1 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
      }
    }
  }
}
  protocols {
    mvrp {
      interface xe-0/1/0.0;
      interface xe-0/1/1.0;
    }
```

Verification

To confirm that the configuration is updating VLAN membership, perform these tasks:

- Verifying That MVRP Is Enabled on Access Switch A on page 1124
- Verifying That MVRP Is Updating VLAN Membership on Access Switch A on page 1124
- Verifying That MVRP Is Enabled on Access Switch B on page 1124

Verifying That MVRP Is Enabled on Access Switch A

Purpose Verify that MVRP is enabled on the switch.

Action Show the MVRP configuration:

```
user@Access-Switch-A> show mvrp
MVRP configuration
MVRP status                  : Enabled
MVRP dynamic VLAN creation   : Enabled

MVRP timers (ms):
Interface            Join   Leave       LeaveAll
--------------       -----  --------    -----------
all                   200   1000        10000
xe-0/1/1.0            200   1000        10000

Interface       Status        Registration Mode
--------------  --------      -----------------
all             Disabled      Normal
xe-0/1/1.0      Enabled       Normal
```

Meaning The results show that MVRP is enabled on the trunk interface of Switch A and that the default timers are used.

Verifying That MVRP Is Updating VLAN Membership on Access Switch A

Purpose Verify that MVRP is updating VLAN membership by displaying the Ethernet switching interfaces and associated VLANs that are active on Switch A.

Action List Ethernet switching interfaces on the switch:

```
user@Access-Switch-A> show ethernet-switching interfaces
Interface   State    VLAN members           Blocking
ge-0/0/1.0  up       finance                unblocked
ge-0/0/2.0  up       lab                    unblocked
ge-0/0/3.0  up       sales                  unblocked
xe-0/1/1.0  up       finance                unblocked
                     lab                    unblocked
```

Meaning MVRP has automatically added **finance** and **lab** as VLAN members on the trunk interface because they are being advertised by access Switch B.

Verifying That MVRP Is Enabled on Access Switch B

Purpose Verify that MVRP is enabled on the switch.

Action Show the MVRP configuration:

```
user@Access-Switch-B> show mvrp

MVRP configuration
```

```
MVRP status                 : Enabled
MVRP dynamic VLAN creation  : Enabled

MVRP timers (ms):
Interface         Join   Leave       LeaveAll
--------------    -----  --------    -----------
all                200   1000       10000
xe-0/1/0.0         200   1000       10000

Interface       Status        Registration Mode
--------------  --------      -----------------
all             Disabled      Normal
xe-0/1/0.0      Enabled       Normal
```

Meaning The results show that MVRP is enabled on the trunk interface of Switch B and that the default timers are used.

Verifying That MVRP Is Updating VLAN Membership on Access Switch B

Purpose Verify that MVRP is updating VLAN membership by displaying the Ethernet switching interfaces and associated VLANs that are active on Switch B.

Action List Ethernet switching interfaces on the switch:

```
user@Access-Switch-B> show ethernet-switching interfaces
Interface   State    VLAN members           Blocking
ge-0/0/0.0  up       finance                unblocked
ge-0/0/1.0  up       lab                    unblocked
xe-0/1/1.0  up       finance                unblocked
                     lab                    unblocked
                     sales                  unblocked
```

Meaning MVRP has automatically added **finance**, **lab**, and **sales** as VLAN members on the trunk interface because they are being advertised by access Switch A.

Verifying That MVRP Is Enabled on Distribution Switch C

Purpose Verify that MVRP is enabled on the switch.

Action Show the MVRP configuration:

```
user@Distribution-Switch-C> show mvrp

MVRP configuration
MVRP status                 : Enabled
MVRP dynamic VLAN creation  : Enabled

MVRP timers (ms):
Interface            Join   Leave       LeaveAll
--------------       -----  --------    -----------
all                   200   1000        10000
xe-0/0/1.0            200   1000        10000
xe-0/1/1.0            200   1000        10000

Interface       Status        Registration Mode
--------------  --------      -----------------
all             Disabled      Normal
```

```
xe-0/0/1.0      Enabled       Normal
xe-0/1/1.0      Enabled       Normal
```

Verifying That MVRP Is Updating VLAN Membership on Distribution Switch C

Purpose Verify that MVRP is updating VLAN membership on distribution Switch C by displaying the Ethernet switching interfaces and associated VLANs on distribution Switch C.

Action List the Ethernet switching interfaces on the switch:

```
user@Distribution-Switch-C> show ethernet-switching interfaces
Interface   State    VLAN members           Blocking
xe-0/1/1.0  up       __mvrp_100__           unblocked
                     __mvrp_200__           unblocked
                     __mvrp_300__           unblocked
xe-0/1/0.0  up       __mvrp_100__           unblocked
                     __mvrp_200__           unblocked
```

List the VLANs that were created dynamically using MVRP on the switch:

```
user@Distribution-Switch-C> show mvrp dynamic-vlan-memberships
VLAN Name                Interfaces
-------------------      ----------------
__mvrp_100__             xe-0/1/1.0
                         xe-0/1/0.0
__mvrp_200__             xe-0/1/1.0
                         xe-0/1/0.0
__mvrp_300__             xe-0/1/1.0
```

Meaning Distribution Switch C has two trunk interfaces. Interface **xe-0/1/1.0** connects distribution Switch C to Access Switch A and is therefore updated to show that it is a member of all the VLANs that are active on Switch A. Any traffic for those VLANs will be passed on from distribution Switch C to Switch A, through interface **xe-0/1/1.0**. Interface **xe-0/1/0.0** connects distribution Switch C to Switch B and is updated to show that it is a member of the two VLANs that are active on Switch B. Thus, distribution Switch C sends traffic for **finance** and **lab** to both Switch A and Switch B. But distribution Switch C sends traffic for **sales** only to Switch A.

Distribution Switch C also has three dynamic VLANs created using MVRP: **mvrp_100**, **mvrp_200**, and **mvrp_300**. The dynamically created VLANs **mvrp_100** and **mvrp_200** are active on interfaces **xe-0/1/1.0** and **xe-0/1/1.0**, and dynamically created VLAN **mvrp_300** is active on interface **xe-0/1/1.0**.

Related Documentation

- Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure) on page 1147
- Understanding Multiple VLAN Registration Protocol (MVRP) on J-EX Series Switches on page 1054

Example: Configuring Layer 2 Protocol Tunneling on J-EX Series Switches

Layer 2 protocol tunneling (L2PT) allows you to send Layer 2 protocol data units (PDUs) across a service provider network and deliver them to J-EX Series switches that are not part of the local broadcast domain. This feature is useful when you want to run Layer 2

protocols on a network that includes switches located at remote sites that are connected across a service provider network.

This example describes how to configure L2PT:

Requirements

This example uses the following hardware and software components:

- Six J-EX Series switches, with three each at two customer sites, with one of the switches at each site designated as the provider edge (PE) device

Overview and Topology

L2PT allows you to send Layer 2 PDUs across a service provider network and deliver them to J-EX Series switches that are not part of the local broadcast domain.

Figure 32 on page 1128 shows a customer network that includes two sites that are connected across a service provider network. Site 1 contains three switches connected in a Layer 2 network, with Switch A designated as a provider edge (PE) device in the service provider network. Site 2 contains a Layer 2 network with a similar topology to that of Site 1, with Switch D designated as a PE device.

Figure 32: L2PT Topology

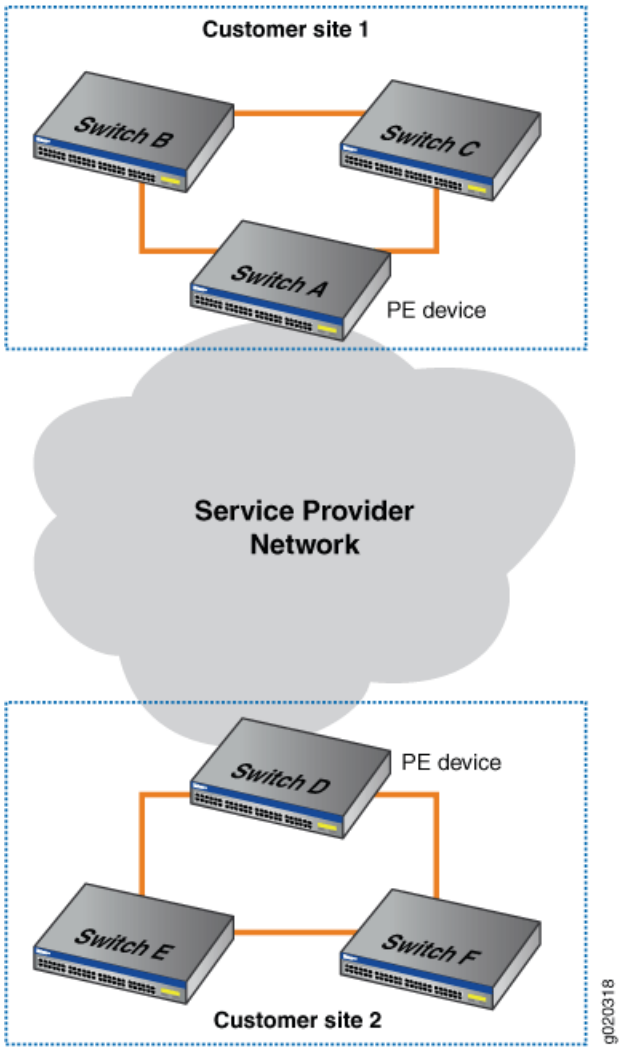

When you enable L2PT on a VLAN, Q-in-Q tunneling is also (and must be) enabled. Q-in-Q tunneling ensures that Switches A, B, C, D, E, and F are part of the same broadcast domain.

This example uses STP as the Layer 2 protocol being tunneled, but you could substitute any of the supported protocols for STP. You can also use the **all** keyword to enable L2PT for all supported Layer 2 protocols.

Tunneled Layer 2 PDUs do not normally arrive at high rate. If the tunneled Layer 2 PDUs do arrive at high rate, there might be a problem in the network. Typically, you would want to shut down the interface that is receiving a high rate of tunneled Layer 2 PDUs so that problem can be isolated. However, if you do not want to completely shut down the interface, you can configure the switch to drop tunneled Layer 2 PDUs that exceed a certain threshold.

The **drop-theshold** configuration statement allows you to specify the maximum number of Layer 2 PDUs of the specified protocol that can be received per second on the interfaces in a specified VLAN before the switch begins dropping the Layer 2 PDUs. The drop threshold must be less than or equal to the shutdown threshold. If the drop threshold is greater than the shutdown threshold and you try to commit the configuration, the commit will fail.

The **shutdown-threshold** configuration statement allows you to specify the maximum number of Layer 2 PDUs of the specified protocol that can be received per second on the interfaces in a specified VLAN before the specified interface is disabled. The shutdown threshold must be greater than or equal to the drop threshold. You can specify a drop threshold without specifying a shutdown threshold, and you can specify a shutdown threshold without specifying a drop threshold. If you do not specify these thresholds, then no thresholds are enforced. As a result, the switch tunnels all Layer 2 PDUs regardless of the speed at which they are received, although the number of packets tunneled per second might be limited by other factors.

In this example, we will configure both a drop threshold and a shutdown threshold to show how this is done.

If L2PT-encapsulated packets are received on an access interface, the switch reacts as it does when there is a loop between the service provider network and the customer network and shuts down (disables) the access interface.

Once an interface is disabled, you must explicitly reenable it using the **clear ethernet-switching layer2-protocol-tunneling error** command or else the interface will remain disabled.

Configuration

To configure L2PT, perform these tasks:

CLI Quick Configuration

To quickly configure L2PT, copy the following commands and paste them into the switch terminal window of each PE device (in Figure 32 on page 1128, Switch A and Switch D are the PE devices):

```
[edit]
set vlans customer-1 dot1q-tunneling
set vlans customer-1 dot1q-tunneling layer2-protocol-tunneling stp
set vlans customer-1 dot1q-tunneling layer2-protocol-tunneling stp drop-threshold 50
set vlans customer-1 dot1q-tunneling layer2-protocol-tunneling stp shutdown-threshold 100
```

Step-by-Step Procedure

To configure L2PT, perform these tasks on each PE device (in Figure 32 on page 1128, Switch A and Switch D are the PE devices):

1. Enable Q-in-Q tunneling on VLAN **customer-1**:

   ```
   [edit]
   user@switch# set vlans customer-1 dot1q-tunneling
   ```

2. Enable L2PT for STP on VLAN **customer-1**:

   ```
   [edit]
   user@switch# set vlans customer-1 dot1q-tunneling layer2-protocol-tunneling stp
   ```

3. Configure the drop threshold as **50**:

```
[edit]
user@switch# set vlans customer-1 dot1q-tunneling layer2–protocol-tunneling stp
drop-threshold 50
```

4. Configure the shutdown threshold as **100**:

```
[edit]
user@switch# set vlans customer-1 dot1q-tunneling layer2–protocol-tunneling stp
shutdown-threshold 100
```

Results Check the results of the configuration:

```
[edit]
user@switch# show vlans customer-1 dot1q-tunneling
layer2-protocol-tunneling {
  stp {
    drop-threshold 50;
    shutdown-threshold 100;
  }
}
```

Verification

To verify that L2PT is working correctly, perform this task:

- Verify That L2PT Is Working Correctly on page 1130

Verify That L2PT Is Working Correctly

Purpose Verify that Q-in-Q tunneling and L2PT are enabled.

Action Check to see that Q-in-Q tunneling and L2PT are enabled on each PE device (Switch A and Switch D are the PE devices):

```
user@switchA> show vlans extensive customer-1
VLAN: customer-1, Created at: Thu Jun 25 05:07:38 2009
802.1Q Tag: 100, Internal index: 4, Admin State: Enabled, Origin: Static
Dot1q Tunneling status: Enabled
Layer2 Protocol Tunneling status: Enabled
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 0 (Active = 0), Untagged  3 (Active = 0)
      ge-0/0/7.0, untagged, access
      ge-0/0/8.0, untagged, access
      ge-0/0/9.0, untagged, access
```

Check to see that L2PT is tunneling STP on VLAN **customer-1** and that **drop-threshold** and **shutdown-threshold** have been configured:

```
user@switchA> show ethernet-switching layer2-protocol-tunneling vlan customer-1

Layer2 Protocol Tunneling VLAN information:
VLAN          Protocol      Drop       Shutdown
                            Threshold  Threshold
customer-1    stp           50         100
```

Check the state of the interfaces on which L2PT has been enabled, including what kind of operation (encapsulation or decapsulation) they are performing:

```
user@switchA> show ethernet-switching layer2-protocol-tunneling interface

Layer2 Protocol Tunneling information:
Interface     Operation       State      Description
ge-0/0/0.0    Encapsulation   Shutdown   Shutdown threshold exceeded
ge-0/0/1.0    Decapsulation   Shutdown   Loop detected
ge-0/0/2.0    Decapsulation   Active
```

Meaning The **show vlans extensive customer-1** command shows that Q-in-Q tunneling and L2PT have been enabled. The **show ethernet-switching layer2-protocol-tunneling vlan customer-1** command shows that L2PT is tunneling the STP protocol on VLAN **customer-1**, the drop threshold is set to **50**, and the shutdown threshold is set to **100**. The **show ethernet-switching layer2-protocol-tunneling interface** command shows the type of operation being performed on each interface, the state of each interface and, if the state is **Shutdown**, the reason why the interface is shut down.

Related Documentation

- Configuring Layer 2 Protocol Tunneling on J-EX Series Switches (CLI Procedure) on page 1150
- Understanding Layer 2 Protocol Tunneling on J-EX Series Switches on page 1056

CHAPTER 59

Configuring Bridging and VLANs

Configuring VLANs for J-EX Series Switches (J-Web Procedure)

You can use the VLAN Configuration page to add a new VLAN or to edit or delete an existing VLAN on a J-EX Series switch.

To access the VLAN Configuration page:

1. Select **Configure** > **Switching** > **VLAN**.

   The VLAN Configuration page displays a list of existing VLANs. If you select a specific VLAN, the specific VLAN details are displayed in the Details section.

NOTE: After you make changes to the configuration in this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See "Using the Commit Options to Commit Configuration Changes (J-Web Procedure)" on page 334s for details about all commit options.

2. Click one:
   - **Add**—creates a VLAN.
   - **Edit**—edits an existing VLAN configuration.
   - **Delete**—deletes an existing VLAN.

NOTE: If you delete a VLAN, the VLAN configuration for all the associated interfaces is also deleted.

When you are adding or editing a VLAN, enter information as described in Table 150 on page 1134.

Table 150: VLAN Configuration Details

Field	Function	Your Action
General tab		
VLAN Name	Specifies a unique name for the VLAN.	Enter a name.
VLAN Id/Range	Specifies the identifier or range for the VLAN.	Select one: • VLAN ID—Type a unique identification number from **1** through **4094**. If no value is specified, it defaults to 0. • VLAN Range—Type a number range to create VLANs with IDs corresponding to the range. For example, the range **2–3** will create two VLANs with the IDs **2** and **3**.
Description	Describes the VLAN.	Enter a brief description for the VLAN.
MAC-Table-Aging-Time	Specifies the maximum time that an entry can remain in the forwarding table before it 'ages out'.	Type the number of seconds from **60** through **1000000**.
Input filter	Specifies the VLAN firewall filter that is applied to incoming packets.	To apply an input firewall filter, select the firewall filter from the list.
Output filter	Specifies the VLAN firewall filter that is applied to outgoing packets.	To apply an output firewall filter, select the firewall filter from the list.
Ports tab		

Table 150: VLAN Configuration Details *(continued)*

Field	Function	Your Action
Ports	Specifies the ports (interfaces) to be associated with this VLAN for data traffic. You can also remove the port association.	Click one: • **Add**—Select the ports from the available list. • **Remove**—Select the port that you do not want associated with the VLAN.
IP address tab		
IPv4 address	Specifies IPv4 address options for the VLAN.	Select **IPv4 address** to enable the IPv4 address options. To configure IPv4: 1. Enter the IP address. 2. Enter the subnet mask—for example, **255.255.255.0**. You can also specify the address prefix. 3. To apply an input firewall filter to an interface, select the firewall filter from the list. 4. To apply an output firewall filter to an interface, select the firewall filter from the list. 5. Click the **ARP/MAC Details** button. Enter the static IP address and MAC address in the window that is displayed.
IPv6 address	Specifies IPv6 address options for the VLAN.	Select **IPv6 address** to enable the IPv6 address options. To configure IPv6: 1. Enter the IP address—for example: **2001:ab8:85a3::8a2e:370:7334**. 2. Specify the subnet mask.
Voip tab		
Ports	Specifies the ports to be associated with this VLAN for voice traffic. You can also remove the port association. NOTE: VoIP is not supported on J-EX8200 switches.	Click one: • **Add**—Select the ports from the available list. • **Remove**—Select the port that you do not want associated with the VLAN.

Related Documentation

- Configuring VLANs for J-EX Series Switches (CLI Procedure) on page 1136
- Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch on page 1063
- Understanding Bridging and VLANs on J-EX Series Switches on page 1041
- Configuring Routed VLAN Interfaces (CLI Procedure) on page 1137

Configuring VLANs for J-EX Series Switches (CLI Procedure)

J-EX Series switches use VLANs to make logical groupings of network nodes with their own broadcast domains. You can use VLANs to limit the traffic flowing across the entire LAN and reduce collisions and packet retransmissions.

For each endpoint on the VLAN, configure the following VLAN parameters on the corresponding interface:

1. Set the description of the VLAN:

```
[edit interfaces interface-name unit 0]
user@switch# set description vlan-description
```

2. Set the unique name of the VLAN:

```
[edit interfaces interface-name unit 0]
user@switch# set family ethernet-switching vlan members vlan-name
```

3. Create the subnet for the VLAN:

```
[edit interfaces]
user@switch# set vlan unit 0 family inet address ip-address
```

4. Configure the VLAN tag ID or VLAN ID range for the VLAN:

```
[edit vlans]
user@switch# set vlan-name vlan-id vlan-id-number
```

   or

```
[edit vlans]
user@switch# set vlan-name vlan-range vlan-id-low-vlan-id-high
```

5. To specify the maximum time that an entry can remain in the forwarding table before it ages out (optional):

```
[edit vlans]
user@switch# set vlan-name mac-table-aging-time time
```

6. To specify a VLAN firewall filter to be applied to incoming or outgoing packets (optional):

```
[edit vlans]
user@switch# set vlan-name filter (input | output) filter-name
```

Related Documentation

- Configuring VLANs for J-EX Series Switches (J-Web Procedure) on page 1133
- Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch on page 1063
- Configuring Routed VLAN Interfaces (CLI Procedure) on page 1137
- Creating a Series of Tagged VLANs (CLI Procedure) on page 1140
- Understanding Bridging and VLANs on J-EX Series Switches on page 1041

Configuring Routed VLAN Interfaces (CLI Procedure)

Routed VLAN interfaces (RVIs) enable the J-EX Series switch to recognize which packets are being sent to local addresses so that they are bridged (switched) whenever possible and are routed only when needed. Whenever packets can be switched instead of routed, several layers of processing are eliminated. Switching also reduces the number of address lookups.

An interface named **vlan** functions as the logical router, on which you can configure a Layer 3 logical interface for each VLAN. For redundancy, an RVI can be combined with implementations of the Virtual Router Redundancy Protocol (VRRP) in both bridging and VPLS environments.

Jumbo frames of up to 9216 bytes are supported on an RVI. To route jumbo data packets on the RVI, you must configure the jumbo MTU size on the member physical interfaces of the RVI and not on the RVI itself (the **vlan** interface). However, for jumbo control packets—for example, to ping the RVI with a packet size of 6000 bytes or more—you must explicitly configure the jumbo MTU size on the interface named **vlan** (the RVI).

CAUTION: Setting or deleting the jumbo MTU size on the RVI (the vlan interface) while the switch is transmitting packets might result in dropped packets.

To configure the routed VLAN interface (RVI):

1. Create a Layer 2 VLAN by assigning it a name (for example, **support**) and a VLAN ID (for example, **111**).

   ```
   [edit]
   user@switch# set vlans support vlan-id 111
   ```

2. Assign an interface (for example, **ge-0/0/18**) to the VLAN (**support**) by naming the VLAN as a trunk member on the logical interface, thereby making the interface part of the VLAN's broadcast domain.

   ```
   [edit]
   user@switch# set interfaces ge-0/0/18 unit 0 family ethernet-switching vlan members
   support
   ```

3. Create a logical Layer 3 RVI (**vlan.111**) on a subnet for the VLAN's broadcast domain.

   ```
   [edit]
   user@switch# set interfaces vlan unit 111 family inet address 111.111.111.1/24
   ```

4. Link the Layer 2 VLAN to the logical Layer 3 interface.

   ```
   [edit]
   user@switch# set vlans support l3-interface vlan.111
   ```

NOTE: Layer 3 interfaces on trunk ports allow the interface to transfer traffic between multiple VLANs. Within a VLAN, traffic is bridged, while across VLANs, traffic is routed.

You can display the configuration settings:

```
user@switch> show interfaces vlan terse
Interface               Admin Link Proto    Local                 Remote
vlan                    up    up
vlan.111                up    up   inet     111.111.111.1/24

user@switch> show vlans
Name            Tag     Interfaces
default
                        None
employee-vlan   20
                        ge-1/0/0.0, ge-1/0/1.0, ge-1/0/2.0
marketing       40
                        ge-1/0/10.0, ge-1/0/20.0, ge-1/0/30.0
support         111
                        ge-0/0/18.0
mgmt
                        bme0.32769, bme0.32771*

user@switch> show ethernet-switching table
Ethernet-switching table: 1 entries, 0 learned
  VLAN              MAC address       Type         Age Interfaces
  support           00:19:e2:50:95:a0 Static         - Router
```

Related Documentation

- Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches on page 1070
- Example: Connecting an Access Switch to a Distribution Switch on page 1078
- Example: Configuring IP Directed Broadcast on a J-EX Series Switch on page 904
- Understanding Bridging and VLANs on J-EX Series Switches on page 1041

Configuring MAC Table Aging (CLI Procedure)

The aging process ensures that the J-EX Series switch tracks only active nodes on the network and that it is able to flush out network nodes that are no longer available.

To manage MAC entries more efficiently, you can configure an entry's aging time, which is the maximum time that an entry can remain in the Ethernet Switching table before it “ages out”.

To configure how long entries remain in the Ethernet Switching table before expiring, using the CLI (here, the VLAN is **employee-vlan**):

```
[edit vlans employee-vlan]
user@switch# set mac-table-aging-time 200
```

Related Documentation

- Understanding Bridging and VLANs on J-EX Series Switches on page 1041
- Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch on page 1063

- Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches on page 1070
- Example: Connecting an Access Switch to a Distribution Switch on page 1078

Configuring the Native VLAN Identifier (CLI Procedure)

J-EX Series switches support receiving and forwarding routed or bridged Ethernet frames with 802.1Q VLAN tags. The logical interface on which untagged packets are to be received must be configured with the same native VLAN ID as that configured on the physical interface.

To configure the native VLAN ID using the CLI:

1. Configure the port mode so that the interface is in multiple VLANs and can multiplex traffic between different VLANs. Trunk interfaces typically connect to other switches and to routers on the LAN. Configure the port mode as **trunk**:

```
[edit interfaces ge-0/0/3 unit 0 family ethernet-switching]
user@switch# set port-mode trunk
```

2. Configure the native VLAN ID:

```
[edit interfaces ge-0/0/3 unit 0 family ethernet-switching]
user@switch# set native-vlan-id 3
```

Related Documentation

- Understanding Bridging and VLANs on J-EX Series Switches on page 1041
- Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches on page 1070
- Example: Connecting an Access Switch to a Distribution Switch on page 1078
- Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch on page 1063

Creating a Series of Tagged VLANs (CLI Procedure)

To identify which VLAN traffic belongs to, all frames on an Ethernet VLAN are identified by a tag, as defined in the IEEE 802.1Q standard. These frames are *tagged* and are encapsulated with 802.1Q tags. For a simple network that has only a single VLAN, all traffic has the same 802.1Q tag.

Instead of configuring VLANS and 802.1Q tags one at a time for a trunk interface, you can configure a VLAN range to create a series of tagged VLANs.

When an Ethernet LAN is divided into VLANs, each VLAN is identified by a unique 802.1Q tag. The tag is applied to all frames so that the network nodes receiving the frames know which VLAN the frames belong to. Trunk ports, which multiplex traffic among a number of VLANs, use the tag to determine the origin of frames and where to forward them.

For example, you could configure the VLAN **employee** and specify a tag range of **10-12**. This creates the following VLANs and tags:

- VLAN **employee-10**, tag **10**
- VLAN **employee-11**, tag **11**
- VLAN **employee-12**, tag **12**

Creating tagged VLANs in a series has the following limitations:

- Layer 3 interfaces do not support this feature.
- Because an access interface can only support one VLAN member, access interfaces also do not support this feature.
- Voice over IP (VoIP) configurations do not support a range of tagged VLANs.

To configure a series of tagged VLANs using the CLI (here, the VLAN is **employee**):

1. Configure the series (here, a VLAN series from 120 through 130):

```
[edit]
user@switch# set vlans employee vlan-range 120-130
```

2. Associate a series of tagged VLANs when you configure an interface in one of two ways:

- Include the name of the series:

```
[edit interfaces]
user@switch# set interfaces ge-0/0/22.0 family ethernet-switching vlan members
employee
```

- Include the VLAN range:

```
[edit interfaces]
user@switch# set interfaces ge-0/0/22.0 family ethernet-switching vlan members
120–130
```

Associating a series of tagged VLANS to an interface by name or by VLAN range have the same result: VLANs **__employee_120__** through **__employee_130__** are created.

NOTE: When a series of VLANs are created using the vlan-range command, the VLAN names are prefixed and suffixed with a double underscore.

Related Documentation

- Verifying That a Series of Tagged VLANs Has Been Created on page 1155
- Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch on page 1063
- Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches on page 1070
- Example: Connecting an Access Switch to a Distribution Switch on page 1078
- Understanding Bridging and VLANs on J-EX Series Switches on page 1041

Configuring Virtual Routing Instances (CLI Procedure)

Use virtual routing and forwarding (VRF) to divide a J-EX Series switch into multiple virtual routing instances. VRF allows you to isolate traffic traversing the network without using multiple devices to segment your network. VRF is supported on all Layer 3 interfaces.

Before you begin, make sure to set up your VLANs. See "Configuring VLANs for J-EX Series Switches (CLI Procedure)" on page 1136 or "Configuring VLANs for J-EX Series Switches (J-Web Procedure)" on page 1133.

To configure virtual routing instances:

1. Create a routing instance:

   ```
   [edit routing-instances]
   user@switch# set routing-instance-name instance-type virtual-router
   ```

   

   NOTE: J-EX Series switches only support the virtual-router instance type.

2. Bind each routing instance to the corresponding physical interfaces:

   ```
   [edit routing-instances]
   user@switch# set routing-instance-name interface interface-name.logical-unit-number
   ```

3. Create the logical interfaces that are bound to the routing instance.

   - To create a logical interface with an IPv4 address:

     ```
     [edit interfaces]
     user@switch# set interface-name unit logical-unit-number family inet address ip-address
     ```

   - To create a logical interface with an IPv6 address:

     ```
     [edit interfaces]
     user@switch# set interface-name unit logical-unit-number family inet6 address
     ipv6–address
     ```

     

     NOTE: Do not create a logical interface using the family ethernet-switching option in this step. Binding an interface using the family ethernet-switching option to a routing instance can cause the interface to shutdown.

4. Enable VLAN tagging on each physical interface that was bound to the routing instance:

   ```
   [edit interfaces]
   user@switch# set interface-name vlan-tagging
   ```

Related Documentation

- Example: Using Virtual Routing Instances to Route Among VLANs on J-EX Series Switches on page 1112
- Verifying That Virtual Routing Instances Are Working on page 1157

- Understanding Virtual Routing Instances on J-EX Series Switches on page 1048

Creating a Private VLAN (CLI Procedure)

The private VLAN (PVLAN) feature on J-EX Series switches allows an administrator to split a broadcast domain into multiple isolated broadcast subdomains, essentially putting a VLAN inside a VLAN.

Before you begin, make sure you set up your VLANs. See “Configuring VLANs for J-EX Series Switches (CLI Procedure)” on page 1136 or “Configuring VLANs for J-EX Series Switches (J-Web Procedure)” on page 1133.

NOTE: Configuring a voice over IP (VoIP) VLAN on PVLAN interfaces is not supported.

To configure a private VLAN:

1. Set the primary VLAN to have no local switching:

   NOTE: The primary VLAN must be a tagged VLAN.

   ```
   [edit vlans]
   user@switch# set primary-vlan-name no-local-switching
   ```

2. For each community VLAN, configure access interfaces:

   NOTE: The secondary VLANs must be untagged VLANs.

   ```
   [edit vlans]
   user@switch# set community-vlan-name interface interface-name
   ```

3. For each community VLAN, set the primary VLAN:

   ```
   [edit vlans]
   user@switch# set community-vlan-name primary-vlan primary-vlan-name
   ```

4. For each isolated VLAN, add the interface to the primary VLAN:

   ```
   [edit vlans]
   user@switch# set primary-vlan-name interface interface-name
   ```

Related Documentation

- Example: Configuring a Private VLAN on a J-EX Series Switch on page 1107
- Verifying That a Private VLAN Is Working on page 1159
- Understanding Private VLANs on J-EX Series Switches on page 1047

Configuring Q-in-Q Tunneling (CLI Procedure)

Q-in-Q tunneling allows service providers on Ethernet access networks to segregate or bundle customer traffic into different VLANs by adding another layer of 802.1Q tags. You can configure Q-in-Q tunneling on J-EX Series switches.

Before you begin configuring Q-in-Q tunneling, make sure you set up your VLANs. See "Configuring VLANs for J-EX Series Switches (CLI Procedure)" on page 1136 or "Configuring VLANs for J-EX Series Switches (J-Web Procedure)" on page 1133.

To configure Q-in-Q tunneling:

1. Enable Q-in-Q tunneling on the S-VLAN:

   ```
   [edit vlans]
   user@switch# set s-vlan-name dot1q-tunneling
   ```

2. Set the allowed C-VLANs on the S-VLAN (optional). Here, the C-VLANs are identified by VLAN range:

   ```
   [edit vlans]
   user@switch# set s-vlan-name dot1q-tunneling customer-vlans range
   ```

3. Change the global Ethertype value (optional):

   ```
   [edit]
   user@switch# set ethernet-switching-options dot1q-tunneling ether-type
   ether-type-value
   ```

4. Disable MAC address learning on the S-VLAN (optional):

   ```
   [edit vlans]
   user@switch# set s-vlan-name no-mac-learning
   ```

Related Documentation

- Example: Setting Up Q-in-Q Tunneling on J-EX Series Switches on page 1105
- Verifying That Q-in-Q Tunneling Is Working on page 1158
- Understanding Q-in-Q Tunneling on J-EX Series Switches on page 1051

Configuring GVRP (J-Web Procedure)

As a network expands and the number of clients and VLANs increases, VLAN administration becomes complex, and the task of efficiently configuring VLANs on multiple J-EX Series switches becomes increasingly difficult. To automate VLAN administration, you can enable GARP VLAN Registration Protocol (GVRP) on the network.

GVRP learns VLANs on a particular 802.1Q trunk port and adds the corresponding trunk interface to the VLAN if the advertised VLAN is preconfigured or already exists on the switch. For example, a VLAN named "sales" is advertised to trunk interface 1 on the GVRP-enabled switch. The switch adds trunk interface 1 to the sales VLAN if the sales VLAN already exists on the switch.

As individual interfaces become active and send requests to join a VLAN, the VLAN configuration is updated and propagated among the switches. Limiting the VLAN configuration to active participants reduces the network overhead. GVRP also provides the benefit of pruning VLANs to limit the scope of broadcast, unknown unicast, and multicast (BUM) traffic to interested network devices only.

To configure GVRP using the J-Web interface:

1. Select **Configure > Switching > GVRP**. Interfaces on which GVRP has been enabled are listed.

   

   NOTE: After you make changes to the configuration in this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See "Using the Commit Options to Commit Configuration Changes (J-Web Procedure)" on page 334 for details about all commit options.

2. To enable GVRP on an interface, click **Add**. Click the arrow key to move the interface from the **Interface Out of GVRP** list to the **Interface under GVRP** list, and click **OK**.
3. To modify GVRP timers, click **Global Settings**. For modifying GVRP Timer settings for the interface, enter information as described in Table 151 on page 1145.
4. Click **OK** to apply changes to the configuration or click **Cancel** to cancel without saving changes.

To disable an interface, select the interface and click **Disable Port**.

Table 151: GVRP Timer Settings

Field	Function	Your Action
Join Timer	Specifies the maximum number of milliseconds the interface waits before sending VLAN advertisements.	Type a number.
Leave Timer	Specifies the number of milliseconds an interface waits after receiving a leave message before the interface leaves the VLAN specified in the message.	Type a number.
Leave All Timer	Specifies the interval in milliseconds at which Leave All messages are sent on interfaces. Leave All messages help to maintain current GVRP VLAN membership information in the network.	Type a number.
Disable GVRP	Disables GVRP on all interfaces.	To disable GVRP, select the check box. To enable GVRP, clear the check box.

Related Documentation

- Example: Configure Automatic VLAN Administration Using GVRP on page 1087
- Monitoring GVRP on page 1161

Configuring Redundant Trunk Groups (J-Web Procedure)

A redundant trunk link provides a simple solution for network recovery when a trunk interface goes down. Traffic is routed to another trunk interface, keeping network convergence time to a minimum. You can configure redundant trunk groups (RTGs) with a primary link and a secondary link on trunk interfaces, or configure dynamic selection of the active interface. If the primary link fails, the secondary link automatically takes over without waiting for normal STP convergence. An RTG can be created only if the following conditions are satisfied:

- A minimum of two trunk interfaces that are not part of any RTG are available.
- All the selected trunk interfaces to be added to the RTG have the same VLAN configuration.
- The selected trunk interfaces are not part of a spanning-tree configuration.

To configure an RTG using the J-Web interface:

1. Select **Configure** > **Switching** > **RTG**.

   The RTG Configuration page displays a list of existing RTGs. If you select a specific RTG, the details of the selected RTG are displayed in the Details of group section.

   

   NOTE: After you make changes to the configuration in this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select Commit Options > Commit. See "Using the Commit Options to Commit Configuration Changes (J-Web Procedure)" on page 334 for details about all commit options.

2. Click one:
   - **Add**—Creates an RTG.
   - **Edit**—Modifies an RTG.
   - **Delete**—Deletes an RTG.

   When you are adding or editing an RTG, enter information as described in Table 152 on page 1146.

3. Click **OK** to apply changes to the configuration or click **Cancel** to cancel without saving changes.

Table 152: RTG Configuration Fields

Field	Function	Your Action
Group Name	Specifies a unique name for the RTG.	Enter a name.
Member Interface 1	Specifies a logical interface containing multiple trunk interfaces.	Select a trunk interface from the list.

Table 152: RTG Configuration Fields *(continued)*

Field	Function	Your Action
Member Interface 2	Specifies a trunk interface containing multiple VLANs.	Select a trunk interface from the list.
Select Primary Interface	Enables you to specify one of the interfaces in the RTG as the primary link. The interface without this option is the secondary link in the RTG.	1. Select the option button. 2. Select the primary interface.
Dynamically select my active interface	Specifies that the system dynamically selects the active interface.	Select the option button.

Related Documentation

- Example: Configuring Redundant Trunk Links for Faster Recovery on page 1101
- Understanding Redundant Trunk Links on J-EX Series Switches on page 1049

Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure)

Multiple VLAN Registration Protocol (MVRP) is used to manage dynamic VLAN registration in a LAN. You can use MVRP on J-EX Series switches.

MVRP is disabled by default on J-EX Series switches.

To enable MVRP or set MVRP options, follow these instructions:

- Enabling MVRP on page 1147
- Disabling MVRP on page 1147
- Disabling Dynamic VLANs on page 1148
- Configuring Timer Values on page 1148
- Configuring MVRP Registration Mode on page 1149

Enabling MVRP

MVRP can only be enabled on trunk interfaces.

To enable MVRP on all trunk interfaces on the switch:

```
[edit protocols mvrp]
user@switch# set interface all
```

To enable MVRP on a specific trunk interface:

```
[edit protocols mvrp]
user@switch# set interface xe-0/0/1.0
```

Disabling MVRP

MVRP is disabled by default. You only need to perform this procedure if you have previously enabled MVRP.

To disable MVRP on all trunk interfaces on the switch:

```
[edit protocols mvrp]
user@switch# set disable
```

To disable MVRP on a specific trunk interface:

```
[edit protocols mvrp]
user@switch# set disable interface xe-0/0/1.0
```

Disabling Dynamic VLANs

Dynamic VLANs can be created on interfaces participating in MVRP by default. Dynamic VLANs are VLANs created on one switch that are propagated to other switches dynamically; in this case, using MVRP.

Dynamic VLAN creation through MVRP cannot be disabled per switch interface. To disable dynamic VLAN creation for interfaces participating in MVRP, you must disable it for all interfaces on the switch.

To disable dynamic VLAN creation:

```
[edit protocols mvrp]
user@switch# set no-dynamic-vlan
```

Configuring Timer Values

The timers in MVRP define the amount of time an interface waits to join or leave MVRP or to send or process the MVRP information for the switch after receiving an MVRP PDU. The join timer controls the amount of time the switch waits to accept a registration request, the leave timer controls the period of time that the switch waits in the Leave state before changing to the unregistered state, and the leaveall timer controls the frequency with which the LeaveAll messages are communicated.

The default MVRP timer values are 200 ms for the join timer, 1000 ms for the leave timer, and 10000 ms for the leaveall timer.

BEST PRACTICE: Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.

To set the join timer for all interfaces on the switch:

```
[edit protocols mvrp]
user@switch# set interface all join-timer 300
```

To set the join timer for a specific interface:

```
[edit protocols mvrp]
user@switch# set interface xe-0/0/1.0 300
```

To set the leave timer for all interfaces on the switch:

```
[edit protocols mvrp]
user@switch# set interface all leave-timer 1200
```

To set the leave timer for a specific interface:

```
[edit protocols mvrp]
user@switch# set interface xe-0/0/1.0 leave-timer 1200
```

To set the leaveall timer for all interfaces on the switch:

```
[edit protocols mvrp]
user@switch# set interface all leaveall-timer 12000
```

To set the leaveall timer for a specific interface:

```
[edit protocols mvrp]
user@switch# set interface xe-0/0/1.0 leaveall-timer 12000
```

Configuring MVRP Registration Mode

The default MVRP registration mode for any interface participating in MVRP is normal. An interface in normal registration mode participates in MVRP when MVRP is enabled on the switch.

An interface in forbidden registration mode does not participate in MVRP even if MVRP is enabled on the switch.

To set all interfaces to forbidden registration mode:

```
[edit protocols mvrp]
user@switch# set interface all registration forbidden
```

To set one interface to forbidden registration mode:

```
[edit protocols mvrp]
user@switch# set interface xe-0/0/1.0 registration forbidden
```

To set all interfaces to normal registration mode:

```
[edit protocols mvrp]
user@switch# set interface all registration normal
```

To set one interface to normal registration mode:

```
[edit protocols mvrp]
user@switch# set interface xe-0/0/1.0 registration normal
```

Related Documentation

- Example: Configuring Automatic VLAN Administration Using MVRP on J-EX Series Switches on page 1115
- Verifying That MVRP Is Working Correctly on page 1162

Configuring Layer 2 Protocol Tunneling on J-EX Series Switches (CLI Procedure)

Layer 2 protocol tunneling (L2PT) allows you to send Layer 2 protocol data units (PDUs) across a service provider network and deliver them to J-EX Series switches at a remote location. This feature is useful when you have a network that includes remote sites that are connected across a service provider network and you want to run Layer 2 protocols on switches connected across the service provider network.

Tunneled Layer 2 PDUs do not normally arrive at high rate. If the tunneled Layer 2 PDUs do arrive at high rate, there might be a problem in the network. Typically, you would want to shut down the interface that is receiving a high rate of tunneled Layer 2 PDUs so that the problem can be isolated. You do so using the **shutdown-threshold** statement. However, if you do not want to completely shut down the interface, you can configure the switch to drop tunneled Layer 2 PDUs that exceed a certain threshold using the **drop-threshold** statement.

There are no default settings for **drop-threshold** and **shutdown-threshold**. If you do not specify these thresholds, then no thresholds are enforced. As a result, the switch tunnels all Layer 2 PDUs regardless of the speed at which they are received, although the number of packets tunneled per second might be limited by other factors.

You can specify a drop threshold value without specifying a shutdown threshold value, and you can specify a shutdown threshold value without specifying a drop threshold value. If you specify both threshold values, then the drop threshold value must be less than or equal to the shutdown threshold value. If the drop threshold value is greater than the shutdown threshold value and you try to commit the configuration, the commit will fail.

NOTE: If the switch receives untagged Layer 2 control PDUs to be tunnelled, then you must configure the switch to map untagged (native) packets to an L2PT-enabled VLAN. Otherwise, the untagged Layer 2 control PDU packets are discarded. For more information, see “Understanding Q-in-Q Tunneling on J-EX Series Switches” on page 1051 and “Configuring Q-in-Q Tunneling (CLI Procedure)” on page 1144.

To configure L2PT on a J-EX Series switch:

1. Because L2PT operates under the Q-in-Q tunneling configuration, you must enable Q-in-Q tunneling before you can configure L2PT. Enable Q-in-Q tunneling on VLAN **customer-1**:

   ```
   [edit]
   user@switch# set vlans customer-1 dot1q-tunneling
   ```

2. Enable L2PT for the Layer 2 protocol you want to tunnel, on the VLAN:

   - To enable L2PT for a specific protocol (here, STP):

     ```
     [edit]
     user@switch# set vlans customer-1 dot1q-tunneling layer2-protocol-tunneling stp
     ```

- To enable L2PT for all supported protocols:

```
[edit]
user@switch# set vlans customer-1 dot1q-tunneling layer2-protocol-tunneling all
```

3. (Optional) Configure the drop threshold:

> NOTE: If you also configure the shutdown threshold, ensure that you configure the drop threshold value to be less than or equal to the shutdown threshold value. If the drop threshold value is greater than the shutdown threshold value and you to try to commit the configuration changes, the commit will fail.

```
[edit]
user@switch# set vlans customer-1 dot1q-tunneling layer2–protocol-tunneling stp
drop-threshold 50
```

4. (Optional) Configure the shutdown threshold:

> NOTE: If you also configure the drop threshold, ensure that you configure the shutdown threshold value to be greater than or equal to the drop threshold value. If the shutdown threshold value is less than the drop threshold value and you to try to commit the configuration changes, the commit will fail.

```
[edit]
user@switch# set vlans customer-1 dot1q-tunneling layer2–protocol-tunneling stp
shutdown-threshold 100
```

> NOTE: Once an interface is disabled, you must explicitly reenable it using the clear ethernet-switching layer2-protocol-tunneling error command. Otherwise, the interface remains disabled.

Related Documentation

- Example: Configuring Layer 2 Protocol Tunneling on J-EX Series Switches on page 1126
- Understanding Layer 2 Protocol Tunneling on J-EX Series Switches on page 1056

Configuring MAC Notification (CLI Procedure)

When a switch learns or unlearns a MAC address, SNMP notifications can be sent to the network management system at regular intervals to record the addition or removal of the MAC address. This process is known as MAC notification.

The MAC notification interval defines how often Simple Network Management Protocol (SNMP) notifications logging the addition or removal of MAC addresses on the switch are sent to the network management system.

MAC notification is disabled by default. When MAC notification is enabled, the default MAC notification interval is 30 seconds.

To enable or disable MAC notification, or to set the MAC notification interval, perform these tasks:

- Enabling MAC Notification on page 1152
- Disabling MAC Notification on page 1152
- Setting the MAC Notification Interval on page 1152

Enabling MAC Notification

MAC notification is disabled by default. You need to perform this procedure to enable MAC notification.

To enable MAC notification on the switch with the default MAC notification interval of 30 seconds:

```
[edit ethernet-switching-options]
user@switch# set mac-notification
```

To enable MAC notification on the switch with any other MAC notification interval (here, the MAC notification interval is set to 60 seconds):

```
[edit ethernet-switching-options]
user@switch# set mac-notification notification-interval 60
```

Disabling MAC Notification

MAC notification is disabled by default. Perform this procedure only if MAC notification was previously enabled on your switch.

To disable MAC notification on the switch:

```
[edit ethernet-switching-options]
user@switch# delete mac-notification
```

Setting the MAC Notification Interval

The default MAC notification interval is 30 seconds. The procedure to change the MAC notification interval to a different interval is identical to the procedure to enable MAC notification on the switch with a nondefault value for the MAC notification interval.

To set the MAC notification interval on the switch (here, the MAC notification interval is set to 5 seconds):

```
[edit ethernet-switching-options]
user@switch# set mac-notification notification-interval 5
```

Related Documentation

- Verifying that MAC Notification Is Working Properly on page 1163

Configuring Proxy ARP (CLI Procedure)

You can configure proxy Address Resolution Protocol (ARP) on your J-EX Series switch to enable the switch to respond to ARP queries for network addresses by offering its own media access control (MAC) address. With proxy ARP enabled, the switch captures and routes traffic to the intended destination.

To configure proxy ARP on a single interface:

```
[edit interfaces]
user@switch# set ge-0/0/3 unit 0 proxy-arp restricted
```

BEST PRACTICE: We recommend that you configure proxy ARP in restricted mode. In restricted mode, the switch is not a proxy if the source and target IP addresses are on the same subnet. If you use unrestricted mode, disable gratuitous ARP requests on the interface to avoid the situation of the switch's response to a gratuitous ARP request appearing to the host to be an indication of an IP conflict:

To configure proxy ARP on a routed VLAN interface (RVI):

```
[edit interfaces]
user@switch# set vlan unit 100 proxy-arp restricted
```

Related Documentation

- Example: Configuring Proxy ARP on a J-EX Series Switch on page 2621
- Verifying That Proxy ARP Is Working Correctly on page 1164
- Configuring Routed VLAN Interfaces (CLI Procedure) on page 1137

CHAPTER 60

Verifying Bridging and VLAN Configuration

Verifying That a Series of Tagged VLANs Has Been Created

Purpose Verify that a series of tagged VLANs is created on the switch.

Action Display the VLANs in the ascending order of their VLAN ID:

```
user@switch> show vlans sort-by tag

Name              Tag         Interfaces
__employee_120__  120
                              ge-0/0/22.0*
__employee_121__  121
                              ge-0/0/22.0*
__employee_122__  122
                              ge-0/0/22.0*
__employee_123__  123
                              ge-0/0/22.0*
__employee_124__  124
                              ge-0/0/22.0*
__employee_125__  125
                              ge-0/0/22.0*
__employee_126__  126
                              ge-0/0/22.0*
__employee_127__  127
                              ge-0/0/22.0*
__employee_128__  128
                              ge-0/0/22.0*
__employee_129__  129
                              ge-0/0/22.0*
```

```
__employee_130__  130
                              ge-0/0/22.0*
```

Display the VLANs by the alphabetical order of the VLAN name:

```
user@switch>  show vlans sort-by name

Name             Tag   Interfaces

__employee_120__ 120
                       ge-0/0/22.0*
__employee_121__ 121
                       ge-0/0/22.0*
__employee_122__ 122
                       ge-0/0/22.0*
__employee_123__ 123
                       ge-0/0/22.0*
__employee_124__ 124
                       ge-0/0/22.0*
__employee_125__ 125
                       ge-0/0/22.0*
__employee_126__ 126
                       ge-0/0/22.0*
__employee_127__ 127
                       ge-0/0/22.0*
__employee_128__ 128
                       ge-0/0/22.0*
__employee_129__ 129
                       ge-0/0/22.0*
__employee_130__ 130
                       ge-0/0/22.0*
```

Display the VLANs by specifying the VLAN-range name (here, the VLAN-range name is **employee**):

```
user@switch>  show vlans employee

Name             Tag   Interfaces

__employee_120__ 120
                       ge-0/0/22.0*
__employee_121__ 121
                       ge-0/0/22.0*
__employee_122__ 122
                       ge-0/0/22.0*
__employee_123__ 123
                       ge-0/0/22.0*
__employee_124__ 124
                       ge-0/0/22.0*
__employee_125__ 125
                       ge-0/0/22.0*
__employee_126__ 126
                       ge-0/0/22.0*
__employee_127__ 127
                       ge-0/0/22.0*
__employee_128__ 128
                       ge-0/0/22.0*
__employee_129__ 129
                       ge-0/0/22.0*
```

```
__employee_130__ 130
                        ge-0/0/22.0*
```

Meaning The sample output shows the VLANs configured on the switch. The series of tagged VLANs is displayed: **__employee_120__** through **__employee_130__**. Each of the tagged VLANs is configured on the trunk interface **ge-0/0/22.0**. The asterisk (*****) beside the interface name indicates that the interface is **UP**.

When a series of VLANs is created using the **vlan-range** statement, the VLAN names are prefixed and suffixed with a double underscore.

Related Documentation

- Creating a Series of Tagged VLANs (CLI Procedure) on page 1140

Verifying That Virtual Routing Instances Are Working

Purpose After creating a virtual routing instance, make sure it is set up properly.

Action

1. Use the **show route instance** command to list all of the routing instances and their properties:

```
user@switch> show route instance

Instance             Type
         Primary RIB
Active/holddown/hidden
master               forwarding
         inet.0                                          3/0/0

__juniper_private1__ forwarding
         __juniper_private1__.inet.0                     1/0/3

__juniper_private2__ forwarding

instance1            forwarding

r1                   virtual-router
         r1.inet.0                                       1/0/0

r2                   virtual-router
         r2.inet.0                                       1/0/0
```

2. Use the **show route forwarding-table** command to view the forwarding table information for each routing instance:

```
user@switch> show route forwarding-table

Routing table: r1.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct   539     2
0.0.0.0/32         perm     0                    dscd   537     1
103.1.1.0/24       ifdn     0                    rslv   579     1
ge-0/0/3.0
103.1.1.0/32       iddn     0 103.1.1.0          recv   577     1
ge-0/0/3.0
103.1.1.1/32       user     0                    rjct   539     2
103.1.1.1/32       intf     0 103.1.1.1          locl   578     2
```

```
103.1.1.1/32       iddn     0 103.1.1.1          locl   578     2
103.1.1.255/32     iddn     0 103.1.1.255        bcst   576     1
ge-0/0/3.0
224.0.0.0/4        perm     0                    mdsc   538     1
224.0.0.1/32       perm     0 224.0.0.1          mcst   534     1
255.255.255.255/32 perm     0                    bcst   535     1
```

Meaning The output confirms that the virtual routing instances are created and the links are up and displays the routing table information.

Related Documentation

- Configuring Virtual Routing Instances (CLI Procedure) on page 1142
- Example: Using Virtual Routing Instances to Route Among VLANs on J-EX Series Switches on page 1112

Verifying That Q-in-Q Tunneling Is Working

Purpose After creating a Q-in-Q VLAN, verify that it is set up properly.

Action

1. Use the **show configuration vlans** command to determine if you successfully created the primary and secondary VLAN configurations:

```
user@switch> show configuration vlans

svlan {
   vlan-id 300;
   dot1q-tunneling {
      customer-vlans [ 101–200 ];
   }
}
```

2. Use the **show vlans** command to view VLAN information and link status:

```
user@switch> show vlans s-vlan-name extensive

VLAN: svlan, Created at: Thu Oct 23 16:53:20 2008
802.1Q Tag: 300, Internal index: 2, Admin State: Enabled, Origin: Static
Dot1q Tunneling Status: Enabled
Customer VLAN ranges:
                     101–200
Protocol: Port Mode
Number of interfaces: Tagged 1 (Active = 0), Untagged  1 (Active = 0)
      ge-0/0/1, tagged, trunk
      ge-0/0/2, untagged, access
```

Meaning The output confirms that Q-in-Q tunnling is enabled and that the VLAN is tagged, and lists the customer VLANs that are associated with the tagged VLAN.

Related Documentation

- Configuring Q-in-Q Tunneling (CLI Procedure) on page 1144
- Example: Setting Up Q-in-Q Tunneling on J-EX Series Switches on page 1105

Verifying That a Private VLAN Is Working

Purpose After creating and configuring private VLANs, verify they are set up properly.

Action

1. Use the **show configuration vlans** command to determine if you successfully created the primary and secondary VLAN configurations:

```
user@switch> show configuration vlans

community1 {
    interface {
        interface a;
        interface b;
    }
    primary-vlan pvlan;
}
community2 {
    interface {
        interface d;
        interface e;
    }
    primary-vlan pvlan;
}
pvlan {
    vlan-id 1000;
    interface {
        isolated1;
        isolated2;
        trunk1;
        trunk2;
    }
    no-local-switching;
}
```

2. Use the **show vlans** command to view VLAN information and link status:

```
user@switch> show vlans pvlan extensive

VLAN: pvlan, Created at: time
802.1Q Tag: vlan-id, Internal index: index-number, Admin State: Enabled,
 Origin: Static
Private VLAN Mode: Primary
Protocol: Port Mode
Number of interfaces: Tagged 2 (Active = 0), Untagged  6 (Active = 0)
      trunk1, tagged, trunk
      interface a, untagged, access
      interface b, untagged, access
      interface c, untagged, access
      interface d, untagged, access
      interface e, untagged, access
      interface f, untagged, access
      trunk2, tagged, trunk
Secondary VLANs: Isolated 2, Community  2
  Isolated VLANs :
      __pvlan_pvlan_isolated1__
      __pvlan_pvlan_isolated2__
  Community VLANs :
      community1
      community2
```

3. Use the **show ethernet-switching table vlan** command to view logs for MAC learning on the VLANs:

```
user@switch> vlan pvlan extensive

pvlan, *
  Interface(s): trunk1
  Interface(s): interface a
  Interface(s): interface b
  Interface(s): interface c
  Interface(s): interface d
  Interface(s): interface e
  Interface(s): interface f
  Interface(s): trunk2
  Type: Flood
  Nexthop index: 1344
```

Meaning The output shows that the primary and secondary VLANs were created and associated and displays MAC learning information.

Related Documentation

- Creating a Private VLAN (CLI Procedure) on page 1143
- Example: Configuring a Private VLAN on a J-EX Series Switch on page 1107

Monitoring Ethernet Switching

Purpose Use the monitoring feature to view details that the J-EX Series switch maintains in its Ethernet switching table. These are details about the nodes on the LAN such as VLAN name, VLAN ID, member interfaces, MAC addresses, and so on.

Action To display Ethernet switching details in the J-Web interface, select **Monitor > Switching > Ethernet Switching**.

To view Ethernet switching details in the CLI, enter the following commands:

- **show ethernet-switching table**
- **show vlans**
- **show ethernet-switching interfaces**

Meaning Table 153 on page 1160 summarizes the Ethernet switching output fields.

Table 153: Ethernet Switching Output Fields

Field	Value
Ethernet Switching Table Information	
MAC Table Count	The number of entries added to the Ethernet switching table.
MAC Table Learned	The number of dynamically learned MAC addresses in the Ethernet switching table.
Ethernet Switching Table Information	

Table 153: Ethernet Switching Output Fields *(continued)*

Field	Value
VLAN	The VLAN name.
MAC Address	The MAC address associated with the VLAN. If a VLAN range has been configured for a VLAN, the output displays the MAC addresses for the entire series of VLANs that were created with that name.
Type	The type of MAC address. Values are: • **static**—The MAC address is manually created. • **learn**—The MAC address is learned dynamically from a packet's source MAC address. • **flood**—The MAC address is unknown and flooded to all members.
Age	The time remaining before the entry ages out and is removed from the Ethernet switching table.
Interfaces	The associated interfaces.
MAC Learning Log	
VLAN-Name	The VLAN name.
MAC Address	The learned MAC address associated with the VLAN ID.
Time	Timestamp for the time at which when the MAC address was added or deleted from the MAC learning log.
State	Operating state of the interface. Values are **Up** and **Down**.

Related Documentation

- Configuring MAC Table Aging (CLI Procedure) on page 1138
- Understanding Bridging and VLANs on J-EX Series Switches on page 1041

Monitoring GVRP

Purpose Use the monitoring feature to view information about the GVRP configuration on the J-EX Series switch.

Action To monitor GVRP in the J-Web interface, select **Monitor > Switching > GVRP**.

To monitor GVRP in the CLI, enter the following command:

- **show gvrp**

Meaning Table 154 on page 1162 summarizes the GVRP output fields.

Table 154: Summary of GVRP Output Fields

Field	Value
Global GVRP Configuration	
GVRP Status	Displays whether GVRP is enabled or disabled.
GVRP Timers	• Join—The number of milliseconds the interfaces must wait before sending VLAN advertisements. • Leave—The number of milliseconds an interface must wait after receiving a Leave message to remove the interface from the VLAN specified in the message. • Leave All—The interval in milliseconds at which Leave All messages are sent on interfaces. Leave All messages maintain current GVRP VLAN membership information in the network.
GVRP Interface Details	
Interface Name	The interface on which GVRP is configured.
Protocol Status	Displays whether GVRP is enabled or disabled on the interface.

Related Documentation

- Configuring GVRP (J-Web Procedure) on page 1144
- Example: Configure Automatic VLAN Administration Using GVRP on page 1087

Verifying That MVRP Is Working Correctly

Purpose After configuring your J-EX Series switch to participate in MVRP, verify that the configuration is properly set and that MVRP messages are being sent and received on your switch.

Action

1. Confirm that MVRP is enabled on your switch.

```
user@switch> show mvrp


Global MVRP configuration
  MVRP status               : Enabled
  MVRP dynamic vlan creation: Enabled
  MVRP Timers (ms):
   Interface        Join   Leave   LeaveAll
   --------------   ----   -----   --------
   all              200    600     10000
   xe-0/1/1.0       200    600     10000

  Interface based configuration:
   Interface        Status     Registration   Dynamic VLAN Creation
   --------------   --------   ------------   ---------------------
   all          Disabled   Fixed          Enabled
   xe-0/1/1.0       Enabled    Normal         Enabled
```

2. Confirm that MVRP messages are being sent and received on your switch.

```
user@switch> show mvrp statistics interface xe-0/1/1.0
```

```
MVRP statistics
  MRPDU received                : 3342
  Invalid PDU received          : 0
  New received                  : 2
  Join Empty received           : 1116
  Join In received              : 2219
  Empty received                : 2
  In received                   : 2
  Leave received                : 1
  LeaveAll received             : 1117
  MRPDU transmitted             : 3280
  MRPDU transmit failures       : 0
  New transmitted               : 0
  Join Empty transmitted        : 1114
  Join In transmitted           : 2163
  Empty transmitted             : 1
  In transmitted                : 1
  Leave transmitted             : 1
  LeaveAll transmitted          : 1111
```

Meaning The output of **show mvrp** shows that interface **xe-0/1/1.0** is enabled for MVRP participation as shown in the status in the **Interface based configuration** field.

The output for **show mvrp statistics interface xe-0/1/1.0** confirms that MVRP messages are being transmitted and received on the interface.

Related Documentation

- Example: Configuring Automatic VLAN Administration Using MVRP on J-EX Series Switches on page 1115
- Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure) on page 1147

Verifying That MAC Notification Is Working Properly

Purpose Verify that MAC notification is enabled or disabled, and that the MAC notification interval is set to the specified value.

Action Verify that MAC notification is enabled while also verifying the MAC notification interval setting.

```
user@switch> show ethernet-switching mac-notification
Notification Status: Enabled
Notification Interval: 30
```

Meaning The output in the **Notification Status** field shows that MAC notification is enabled. The output in the **Notification Status** field would display **Disabled** if MAC notification was disabled.

The **Notification Interval** field output shows that the MAC notification interval is set to 30 seconds.

Related Documentation

- Configuring MAC Notification (CLI Procedure) on page 1151

Verifying That Proxy ARP Is Working Correctly

Purpose Verify that the switch is sending proxy ARP messages.

Action List the system statistics for ARP:

```
user@switch> show system statistics arp
arp:
        198319 datagrams received
        45 ARP requests received
        12 ARP replies received
        2 resolution requests received
        2 unrestricted proxy requests
        0 restricted proxy requests
        0 received proxy requests
        0 proxy requests not proxied
        0 restricted-proxy requests not proxied
        0 with bogus interface
        0 with incorrect length
        0 for non-IP protocol
        0 with unsupported op code
        0 with bad protocol address length
        0 with bad hardware address length
        0 with multicast source address
        0 with multicast target address
        0 with my own hardware address
        168705 for an address not on the interface
        0 with a broadcast source address
        0 with source address duplicate to mine
        29555 which were not for me
        0 packets discarded waiting for resolution
        4 packets sent after waiting for resolution
        27 ARP requests sent
        47 ARP replies sent
        0 requests for memory denied
        0 requests dropped on entry
        0 requests dropped during retry
        0 requests dropped due to interface deletion
        0 requests on unnumbered interfaces
        0 new requests on unnumbered interfaces
        0 replies for from unnumbered interfaces
        0 requests on unnumbered interface with non-subnetted donor
        0 replies from unnumbered interface with non-subnetted donor
```

Meaning The statistics show that two proxy ARP requests were received, and the **proxy requests not proxied** field indicates that all the unproxied ARP requests received have been proxied by the switch.

Related Documentation

- Configuring Proxy ARP (CLI Procedure) on page 1153

CHAPTER 61

Troubleshooting Bridging and VLAN Configuration

Troubleshooting Ethernet Switching

Troubleshooting issues for Ethernet switching on J-EX Series switches:

MAC Address in the Switch's Ethernet Switching Table Is Not Updated After a MAC Address Move

Problem Sometimes a MAC address entry in the switch's Ethernet switching table is not updated after the device with that MAC address has been moved from one interface to another on the switch. Typically, the switch does not wait for a MAC address expiration when a MAC move operation occurs. As soon as the switch detects the MAC address on the new interface, it immediately updates the table. Many network devices send a gratuitous ARP packet when switching an IP address from one device to another. The switch updates its ARP cache table after receipt of such gratuitous ARP messages, and then it also updates its Ethernet switching table. However, sometimes silent devices, such as SYSLOG servers or SNMP Trap receivers that receive UDP traffic but do not return acknowledgement (ACK) messages to the traffic source, do not send gratuitous ARP packets when a device moves. If such a move occurs when the system administrator is not available to explicitly clear the affected interfaces by issuing the **clear ethernet-switching table** command, the entry for the moved device in the Ethernet switching table is not updated.

Solution Set up the switch to handle unattended MAC address switchovers.

1. Reduce the system-wide ARP aging timer. (By default, the ARP aging timer is set at 20 minutes. The range of the timer is from 1 through 240 minutes.)

   ```
   [edit system arp]
   user@switch# set aging-timer 3
   ```

2. Set the MAC aging timer to the same value as the ARP timer. (By default, the MAC aging timer is set to 300 seconds. The range is 15 to 1,000,000 seconds.)

```
[edit vlans]
user@switch# set vlans sales mac-table-aging-time 180
```

The ARP entry and the MAC address entry for the moved device expire within the times specified by the aging timer values. After the entries expire, the switch sends a new ARP message to the IP address of the device. The device responds to the ARP, thereby refreshing the entries in the switch's ARP cache table and Ethernet switching table

Related Documentation

- arp on page 171
- mac-table-aging-time on page 1209

CHAPTER 62

Configuration Statements for Bridging and VLANs

[edit ethernet-switching-options] Configuration Statement Hierarchy

```
ethernet-switching-options {
  analyzer {
    name {
      loss-priority priority;
      ratio number;
      input {
        ingress {
          interface (all | interface-name);
          vlan (vlan-id | vlan-name);
        }
        egress {
          interface (all | interface-name);
        }
      output {
        interface interface-name;
        vlan (vlan-id | vlan-name);
      }
    }
  }
  bpdu-block {
    disable-timeout timeout;
    interface (all | [interface-name]);
  }
  dot1q-tunneling {
    ether-type (0x8100 | 0x88a8 | 0x9100);
  }
  interfaces interface-name {
    no-mac-learning;
  }
```

```
mac-notification {
  notification-interval seconds;
}
mac-table-aging-time seconds;
port-error-disable {
  disable-timeout timeout;
}
redundant-trunk-group {
  group-name name {
    interface interface-name <primary>;
  }
}
secure-access-port {
  dhcp-snooping-file {
    location local_pathname | remote_URL;
    timeout seconds;
    write-interval seconds;
  }
  interface (all | interface-name) {
    allowed-mac {
      mac-address-list;
    }
    (dhcp-trusted | no-dhcp-trusted );
    mac-limit limit action action;
    no-allowed-mac-log;
    static-ip  ip-address {
      vlan  vlan-name;
      mac mac-address;
    }
  }
  vlan (all | vlan-name) {
    (arp-inspection | no-arp-inspection );
    dhcp-option82 {
      circuit-id {
        prefix hostname;
        use-interface-description;
        use-vlan-id;
      }
      remote-id {
        prefix hostname | mac | none;
        use-interface-description;
        use-string string;
      }
      vendor-id [string];
    }
    (examine-dhcp | no-examine-dhcp );
    (ip-source-guard | no-ip-source-guard);
    mac-move-limit limit action action;
  }
}
storm-control {
  action-shutdown;
  interface (all | interface-name) {
    bandwidth bandwidth;
    no-broadcast;
    no-unknown-unicast;
```

```
    }
  }
  traceoptions {
    file filename <files number> <no-stamp> <replace> <size size> <world-readable |
      no-world-readable>;
    flag flag <disable>;
  }
  unknown-unicast-forwarding {
    vlan (all | vlan-name) {
      interface interface-name;
    }
  }
  voip {
    interface (all | [interface-name | access-ports]) {
      vlan vlan-name ;
      forwarding-class <assured-forwarding | best-effort | expedited-forwarding |
        network-control>;
    }
  }
}
```

Related Documentation

- Understanding Port Mirroring on J-EX Series Switches on page 3245
- Port Security for J-EX Series Switches Overview on page 2545
- Understanding BPDU Protection for STP, RSTP, and MSTP on J-EX Series Switches on page 1278
- Understanding Redundant Trunk Links on J-EX Series Switches on page 1049
- Understanding Storm Control on J-EX Series Switches on page 2511
- Understanding 802.1X and VoIP on J-EX Series Switches on page 2263
- Understanding Q-in-Q Tunneling on J-EX Series Switches on page 1051
- Understanding Unknown Unicast Forwarding on J-EX Series Switches on page 2512
- Understanding MAC Notification on J-EX Series Switches on page 1060

[edit interfaces] Configuration Statement Hierarchy

```
interfaces {
  aex {
    aggregated-ether-options {
      (flow-control | no-flow-control);
      lacp mode {
        periodic interval;
      }
      link-speed speed;
      minimum-links number;
    }
    description text;
    disable;
    hold-time up milliseconds down milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;
```

```
    traceoptions;
    (traps | no-traps);
    unit logical-unit-number {
      description text;
      disable;
      family family-name {...}
      proxy-arp (restricted | unrestricted);
      (traps | no-traps);
      vlan-id vlan-id-number;
    }
    vlan-tagging;
  }
  fe-fpc/pic/port {
    description text;
    disable;
    mtu bytes;
    no-gratuitous-arp-request;
    speed speed;
    traceoptions;
    (traps | no-traps);
    unit logical-unit-number {
      description text;
      disable;
      family family-name {...}
      proxy-arp (restricted | unrestricted);
      (traps | no-traps);
      vlan-id vlan-id-number;
    }
    vlan-tagging;
  }
  ge-fpc/pic/port {
    description text;
    disable;
    ether-options {
      802.3ad aex {
        lacp {
          force-up;
        }
      }
      (auto-negotiation | no-auto-negotiation);
      (flow-control | no-flow-control);
      link-mode mode;
      speed (auto-negotiation | speed);
    }
    hold-time up milliseconds down milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;
    traceoptions;
    (traps | no-traps);
    unit logical-unit-number {
      description text;
      disable;
      family family-name {...}
      proxy-arp (restricted | unrestricted);
      rpm;
      (traps | no-traps);
```

```
    vlan-id vlan-id-number;
  }
  vlan-tagging;
}
interface-range interface-range name {
  description text;
  disable;
  ether-options {
    802.3ad aex {
      lacp {
        force-up;
      }
    }
    (auto-negotiation | no-auto-negotiation);
    (flow-control| no-flow-control);
    link-mode mode;
    speed (auto-negotiation | speed);
  }
  hold-time up milliseconds down milliseconds;
  member interface-name;
  member-range starting-interface name to ending-interface name;
  mtu bytes;
  unit logical-unit-number {
    description text;
    disable;
    family family-name {...}
    proxy-arp (restricted | unrestricted);
    rpm;
    (traps | no-traps);
    vlan-id vlan-id-number;
  }
}
lo0 {
  description text;
  disable;
  hold-time up milliseconds down milliseconds;
  traceoptions;
  (traps | no-traps);
  unit logical-unit-number {
    description text;
    disable;
    family family-name {...}
    (traps | no-traps);
  }
}
me0 {
  description text;
  disable;
  hold-time up milliseconds down milliseconds;
  no-gratuitous-arp-request;
  traceoptions;
  (traps | no-traps);
  unit logical-unit-number {
    description text;
    disable;
    family family-name {...}
```

```
      (traps | no-traps);
      vlan-id vlan-id-number;
    }
    vlan-tagging;
  }
  vlan {
    description  text;
    disable;
    hold-time up milliseconds down milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;
    traceoptions;
    (traps | no-traps);
    unit logical-unit-number {
      description text;
      disable;
      family family-name {...}
      proxy-arp (restricted | unrestricted);
      (traps | no-traps);
    }
  }
  vme {
    description text;
    disable;
    hold-time up milliseconds down milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;
    traceoptions;
    (traps | no-traps);
    unit logical-unit-number {
      description text;
      disable;
      family family-name {...}
      (traps | no-traps);
      vlan-id vlan-id-number;
    }
    vlan-tagging;
  }
  xe-fpc/pic/port {
    description text;
    disable;
    ether-options {
      802.3ad aex {
        lacp (802.3ad) {
          force-up;
        }
      }
      (auto-negotiation | no-auto-negotiation);
      (flow-control | no-flow-control);
      link-mode mode;
      speed (auto-negotiation | speed);
    }
    hold-time up milliseconds down milliseconds;
    mtu bytes;
    no-gratuitous-arp-request;
    traceoptions;
```

```
    (traps | no-traps);
    unit logical-unit-number {
      description text;
      disable;
      family family-name {...}
      proxy-arp (restricted | unrestricted);
      rpm;
      (traps | no-traps);
      vlan-id vlan-id-number;
    }
    vlan-tagging;
  }
}
```

Related Documentation

- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 919
- Configuring Aggregated Ethernet Interfaces (CLI Procedure) on page 922
- Configuring a Layer 3 Subinterface (CLI Procedure) on page 930
- Configuring Routed VLAN Interfaces (CLI Procedure) on page 1137
- Configuring the Virtual Management Ethernet Interface for Global Management of a Virtual Chassis (CLI Procedure) on page 797
- J-EX Series Switches Interfaces Overview on page 863
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

[edit protocols] Configuration Statement Hierarchy

```
protocols {
  connections {
    remote-interface-switch connection-name {
      interface interface-name.unit-number;
      transmit-lsp label-switched-path;
      receive-lsp label-switched-path;
    }
  }
  dot1x {
    authenticator {
      authentication-profile-name profile-name;
      interface (all | [ interface-names ]) {
        disable;
        guest-vlan ( vlan-id | vlan-name);
        mac-radius <restrict>;
        maximum-requests number;
        no-reauthentication;
        quiet-period seconds;
        reauthentication {
          interval seconds;
        }
        retries number;
        server-fail (deny | permit | use-cache | vlan-id | vlan-name);
        server-reject-vlan (vlan-id | vlan-name);
```

```
        server-timeout seconds;
        supplicant (multiple | single | single-secure);
        supplicant-timeout seconds;
        transmit-period seconds;
      }
      static mac-address {
        interface interface-name;
        vlan-assignment (vlan-id |vlan-name);
      }
  }
  gvrp {
    <enable | disable>;
    interface (all | [interface-name]) {
      disable;
    }
    join-timer millseconds;
    leave-timer milliseconds;
    leaveall-timer milliseconds;
  }
  igmp-snooping {
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>
        <match regex>;
      flag flag (detail | disable | receive | send);
    }
    vlan (vlan-id | vlan-number) {
      data-forwarding {
        source {
          groups group-prefix;
        }
        receiver {
          source-vlans vlan-list;
          install ;
        }
      }
      disable {
        interface interface-name
      }
      immediate-leave;
      interface interface-name {
        group-limit limit;
        multicast-router-interface;
        static {
          group ip-address;
        }
      }
      proxy ;
      query-interval seconds;
      query-last-member-interval seconds;
      query-response-interval seconds;
      robust-count number;
    }
  }
  lldp {
    disable;
    advertisement-interval seconds;
```

```
    hold-multiplier number;
    interface (all | interface-name) {
      disable;
    }
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>
        <match regex>;
      flag flag (detail | disable | receive | send);
    }
  }
  lldp-med {
    disable;
    fast-start number;
    interface (all | interface-name) {
      disable;
      location {
        elin number;
        civic-based {
          what number;
          country-code code;
          ca-type {
            number {
              ca-value value;
            }
          }
        }
      }
    }
  }
  mpls {
    interface ( all | interface-name );
    label-switched-path lsp-name to remote-provider-edge-switch;
    path destination {
      <address | hostname> <strict | loose>
    }
  mstp {
    disable;
    bpdu-block-on-edge;
    bridge-priority priority;
    configuration-name name;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
      disable;
      bpdu-timeout-action {
        block;
        alarm;
      }
      cost  cost;
      edge;
      mode mode;
      no-root-port;
      priority priority;
    }
    max-age seconds;
    max-hops hops;
```

```
    msti msti-id {
      vlan (vlan-id | vlan-name);
      interface interface-name {
        disable;
        cost  cost;
        edge;
        mode mode;
        priority priority;
      }
    }
    revision-level revision-level;
    traceoptions {
      file filename <files number > <size size> <no-stamp | world-readable |
        no-world-readable>;
      flag flag;
    }
  }
  mvrp {
    disable
    interface (all | interface-name) {
      disable;
      join-timer milliseconds;
      leave-timer milliseconds;
      leaveall-timer milliseconds;
      registration (forbidden | normal);
    }
    no-dynamic-vlan;
    traceoptions {
      file filename <files number > <size size> <no-stamp | world-readable |
        no-world-readable>;
      flag flag;
    }
  }
  oam {
    ethernet{
      connectivity-fault-management {
        action-profile profile-name {
          default-actions {
            interface-down;
          }
        }
        linktrace {
          age (30m | 10m | 1m | 30s | 10s);
          path-database-size path-database-size;
        }
        maintenance-domain domain-name {
          level number;
          mip-half-function (none | default |explicit);
          name-format (character-string | none | dns | mac+2oct);
          maintenance-association ma-name {
            continuity-check {
              hold-interval minutes;
              interval (10m | 10s | 1m | 1s| 100ms);
              loss-threshold number;
            }
            mep mep-id {
```

```
                auto-discovery;
                direction down;
                interface interface-name;
                remote-mep mep-id {
                  action-profile profile-name;
                }
              }
            }
          }
        }
        link-fault-management {
          action-profile profile-name;
            action {
              syslog;
              link-down;
            }
            event {
              link-adjacency-loss;
              link-event-rate;
              frame-error count;
              frame-period count;
              frame-period-summary count;
              symbol-period count;
            }
            interface interface-name {
              link-discovery (active | passive);
              pdu-interval interval;
              event-thresholds threshold-value;
              remote-loopback;
              event-thresholds {
                frame-errorcount;
                frame-period count;
                frame-period-summary count;
                symbol-period count;
              }
            }
            negotiation-options {
              allow-remote-loopback;
              no-allow-link-events;
            }
        }
      }
    }
    rstp {
      disable;
      bpdu-block-on-edge;
      bridge-priority priority;
      forward-delay seconds;
      hello-time seconds;
      interface (all | interface-name) {
        disable;
        bpdu-timeout-action {
          block;
          alarm;
        }
        cost  cost;
```

```
        edge;
        mode mode;
        no-root-port;
        priority priority;
      }
      max-age seconds;
      }
      traceoptions {
        file filename <files number > <size size> <no-stamp | world-readable |
          no-world-readable>;
        flag flag;
      }
  }
  sflow {
    agent-id
    collector {
      ip-address;
      udp-port port-number;
    }
    disable;
    interfaces interface-name {
      disable;
      polling-interval seconds;
      sample-rate number;
    }
    polling-interval seconds;
    sample-rate number;
    source-ip
  }
  stp {
    disable;
    bridge-priority priority;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
      disable;
      bpdu-timeout-action {
        block;
        alarm;
      }
      cost  cost;
      edge;
      mode mode;
      no-root-port;
      priority priority;
    }
    max-age seconds;
  }
    traceoptions {
      file filename <files number > <size size> <no-stamp | world-readable |
        no-world-readable>;
      flag flag;
    }
  vstp {
    bpdu-block-on-edge;
    disable;
```

```
        force-version stp;
        vlan (all | vlan-id | vlan-name) {
          bridge-priority priority;
          forward-delay seconds;
          hello-time seconds;
          interface (all | interface-name) {
            bpdu-timeout-action {
              alarm;
              block;
            }
            cost  cost;
            disable;
            edge;
            mode mode;
            no-root-port;
            priority priority;
          }
          max-age seconds;
          traceoptions {
            file filename <files number > <size size> <no-stamp | world-readable |
              no-world-readable>;
            flag flag;
          }
        }
      }
    }
```

Related Documentation

- 802.1X for J-EX Series Switches Overview on page 2253
- Example: Configure Automatic VLAN Administration Using GVRP on page 1087
- Understanding MAC RADIUS Authentication on J-EX Series Switches
- Understanding Server Fail Fallback and 802.1X Authentication on J-EX Series Switches on page 2258
- IGMP Snooping on J-EX Series Switches Overview on page 2047
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261
- Understanding MSTP for J-EX Series Switches on page 1277
- Understanding Multiple VLAN Registration Protocol (MVRP) on J-EX Series Switches on page 1054
- Understanding Ethernet OAM Connectivity Fault Management for a J-EX Series Switch on page 3463
- Understanding Ethernet OAM Link Fault Management for a J-EX Series Switch on page 3427
- Understanding RSTP for J-EX Series Switches on page 1276
- Understanding STP for J-EX Series Switches on page 1275
- Understanding How to Use sFlow Technology for Network Monitoring on a J-EX Series Switch on page 3283
- Understanding VSTP for J-EX Series Switches on page 1281

[edit routing-instances] Configuration Hierarchy

```
routing-instances routing-instance-name {
  instance-type virtual-router
  interface interface-name
}
```

Related Documentation

- Example: Using Virtual Routing Instances to Route Among VLANs on J-EX Series Switches on page 1112
- Configuring Virtual Routing Instances (CLI Procedure) on page 1142

[edit vlans] Configuration Statement Hierarchy

```
vlans {
  vlan-name {
    description text-description;
    dot1q-tunneling {
      customer-vlans (id | native | range);
      layer2-protocol-tunneling all | protocol-name {
        drop-threshold number;
        shutdown-threshold number;
      }
    }
    filter input filter-name;
    filter output filter-name;
    interface interface-name {
      mapping (native (push | swap) | policy | tag (push | swap));
    }
    l3-interface vlan.logical-interface-number;
    mac-limit number;
    mac-table-aging-time seconds;
    no-local-switching;
    no-mac-learning;
    primary-vlan vlan-name;
    vlan-id number;
    vlan-range vlan-id-low-vlan-id-high;
  }
}
```

Related Documentation

- Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch on page 1063
- Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches on page 1070
- Example: Configure Automatic VLAN Administration Using GVRP on page 1087
- Example: Connecting an Access Switch to a Distribution Switch on page 1078
- Example: Setting Up Q-in-Q Tunneling on J-EX Series Switches on page 1105
- Example: Configuring Layer 2 Protocol Tunneling on J-EX Series Switches on page 1126
- Creating a Private VLAN (CLI Procedure) on page 1143
- Understanding Q-in-Q Tunneling on J-EX Series Switches on page 1051

arp

Syntax

```
arp {
  aging-timer minutes;
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set the time interval between ARP updates.

Options

aging-timer *minutes*—Time interval in minutes between ARP updates. In environments where the number of ARP entries to update is high, increasing the time between updates can improve system performance.

Range: 5 to 240 minutes

Default: 20 minutes

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- For more information about ARP updates, see the *Junos OS System Basics Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.

bridge-priority

Syntax bridge-priority *priority*;

Hierarchy Level [edit protocols **mstp**],
[edit protocols **mstp msti** *msti-id*],
[edit protocols **rstp**],
[edit protocols **stp**],
[edit protocols **vstp vlan** *vlan-id*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the bridge priority. The bridge priority determines which bridge is elected as the root bridge. If two bridges have the same path cost to the root bridge, the bridge priority determines which bridge becomes the designated bridge for a LAN segment.

Default 32,768

Options ***priority***—Bridge priority. It can be set only in increments of 4096.

Range: 0 through 61,440

Default: 32,768

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show spanning-tree bridge on page 1398**
- **show spanning-tree interface on page 1407**
- Example: Configuring Network Regions for VLANs with MSTP on J-EX Series Switches on page 1297
- Understanding MSTP for J-EX Series Switches on page 1277
- Understanding VSTP for J-EX Series Switches on page 1281

customer-vlans

Syntax customer-vlans (*id* | native | *range*);

Hierarchy Level [edit **vlans** *vlan-name* **dot1q-tunneling**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Limit the set of accepted C-VLAN tags to a range or to discrete values.

Options ***id***—Numeric identifier for a VLAN.

native—Accepts untagged and priority-tagged packets from access interfaces and assigns the configured S-VLAN to the packet.

range—Range of numeric identifiers for VLANs.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **dot1q-tunneling on page 1185**
- **ether-type on page 1188**
- Example: Setting Up Q-in-Q Tunneling on J-EX Series Switches on page 1105
- Configuring Q-in-Q Tunneling (CLI Procedure) on page 1144
- Understanding Q-in-Q Tunneling on J-EX Series Switches on page 1051

description

Syntax description *text-description*;

Hierarchy Level [edit **vlans** *vlan-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.
Option **text-description** enhanced from supporting up to 128 characters to supporting up to 256 characters in Junos OS Release 10.2 for J-EX Series switches.

Description Provide a textual description of the VLAN. The text has no effect on the operation of the VLAN or switch.

Options ***text-description***—Text to describe the interface. It can contain letters, numbers, and hyphens (-) and can be up to 256 characters long. If the text includes spaces, enclose the entire text in quotation marks.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **show vlans on page 1263**
- Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch on page 1063
- Understanding Bridging and VLANs on J-EX Series Switches on page 1041

disable

Syntax disable;

Hierarchy Level [edit protocols **gvrp**],
[edit protocols **gvrp interface** [*interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

NOTE: GVRP can be enabled only on trunk interfaces.

Description Disable the GVRP configuration on the interface.

Default If you do not configure GVRP, it is disabled. You can use this command to disable a prior configuration of GVRP on a specified interface.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **show gvrp on page 1253**
- Example: Configure Automatic VLAN Administration Using GVRP on page 1087

disable (MVRP)

Syntax disable;

Hierarchy Level [edit protocols **mvrp**],
[edit protocols **mvrp interface**(all | *interface-name*)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable the MVRP configuration on the interface.

Default MVRP is disabled by default.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure) on page 1147

dot1q-tunneling (Ethernet Switching)

Syntax
```
dot1q-tunneling {
  ether-type (0x8100 | 0x88a8 | 0x9100);
}
```

Hierarchy Level [edit **ethernet-switching-options**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.
The remaining statement is explained separately.

Description Set a global value for the Ethertype for Q-in-Q tunneling.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **dot1q-tunneling on page 1186**
- Example: Setting Up Q-in-Q Tunneling on J-EX Series Switches on page 1105
- Configuring Q-in-Q Tunneling (CLI Procedure) on page 1144

dot1q-tunneling (VLANs)

Syntax

```
dot1q-tunneling {
  customer-vlans (id | native | range);
  layer2-protocol-tunneling all | protocol-name {
    drop-threshold number;
    shutdown-threshold number;
  }
}
```

Hierarchy Level [edit vlans *vlan-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable Q-in-Q tunneling on the specified VLAN.

NOTE: The VLAN on which you enable Q-in-Q tunneling must be a tagged VLAN.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **dot1q-tunneling on page 1185**
- Example: Setting Up Q-in-Q Tunneling on J-EX Series Switches on page 1105
- Example: Configuring Layer 2 Protocol Tunneling on J-EX Series Switches on page 1126
- Configuring Q-in-Q Tunneling (CLI Procedure) on page 1144
- Understanding Q-in-Q Tunneling on J-EX Series Switches on page 1051

drop-threshold

Syntax drop-threshold *number*;

Hierarchy Level [edit vlans *vlan-name* dot1q-tunneling layer2-protocol-tunneling all | *protocol-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the maximum number of Layer 2 PDUs of the specified protocol that can be received per second on the interfaces in a specified VLAN before the switch begins dropping the Layer 2 PDUs. The drop threshold value must be less than or equal to the shutdown threshold value.

NOTE: If the drop threshold value is greater than the shutdown threshold value and you try to commit the configuration, the commit will fail.

You can specify a drop threshold value without specifying a shutdown threshold value.

Default No drop threshold is specified.

Options *number*—Maximum number of Layer 2 PDUs of the specified protocol that can be received per second on the interfaces in a specified VLAN before the switch begins dropping the Layer 2 PDUs.

Range: 1 through 1000

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **shutdown-threshold on page 1220**
- Example: Configuring Layer 2 Protocol Tunneling on J-EX Series Switches on page 1126
- Configuring Layer 2 Protocol Tunneling on J-EX Series Switches (CLI Procedure) on page 1150

ether-type

Syntax ether-type (0x8100 | 0x88a8 | 0x9100)

Hierarchy Level [edit **ethernet-switching-options dot1q-tunneling**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a global value for the Ethertype. Only one Ethertype value is supported at a time. The Ethertype value appears in the Ethernet type field of the packet. It specifies the protocol being transported in the Ethernet frame.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **dot1q-tunneling on page 1186**
- Example: Setting Up Q-in-Q Tunneling on J-EX Series Switches on page 1105
- Configuring Q-in-Q Tunneling (CLI Procedure) on page 1144

ethernet-switching-options

Syntax

```
ethernet-switching-options {
  analyzer {
    name {
      loss-priority priority;
      ratio number;
      input {
        ingress {
          interface (all | interface-name);
          vlan (vlan-id | vlan-name);
        }
        egress {
          interface (all | interface-name);
        }
      output {
        interface interface-name;
        vlan (vlan-id | vlan-name);
      }
    }
  }
  bpdu-block {
    disable-timeout timeout;
    interface (all | [interface-name]);
  }
  dot1q-tunneling {
    ether-type (0x8100 | 0x88a8 | 0x9100);
  }
  interfaces interface-name {
    no-mac-learning;
  }
  mac-notification {
    notification-interval seconds;
  }
  mac-table-aging-time seconds;
  port-error-disable {
    disable-timeout timeout;
  }
  redundant-trunk-group {
    group-name name {
      interface interface-name <primary>;
      interface interface-name;
    }
  }
  secure-access-port {
    dhcp-snooping-file {
      location local_pathname | remote_URL;
      timeout seconds;
      write-interval seconds;
    }
    interface (all | interface-name) {
      allowed-mac {
        mac-address-list;
      }
```

```
        (dhcp-trusted | no-dhcp-trusted);
        mac-limit limit action action;
        no-allowed-mac-log;
        static-ip ip-address {
          vlan  vlan-name;
          mac mac-address;
        }
      }
      vlan (all | vlan-name) {
        (arp-inspection | no-arp-inspection);
        dhcp-option82 {
          circuit-id {
            prefix hostname;
            use-interface-description;
            use-vlan-id;
          }
          remote-id {
            prefix hostname | mac | none;
            use-interface-description;
            use-string string;
          }
          vendor-id [string];
        }
        (examine-dhcp | no-examine-dhcp);
        (ip-source-guard | no-ip-source-guard);
        mac-move-limit limit action action;
      }
    }
    storm-control {
      action-shutdown;
      interface (all | interface-name) {
        bandwidth bandwidth;
        no-broadcast;
        no-unknown-unicast;
      }
    }
    traceoptions {
      file filename <files number> <no-stamp> <replace> <size size> <world-readable |
        no-world-readable>;
      flag flag <disable>;
    }
    unknown-unicast-forwarding {
      vlan (all | vlan-name) {
        interface interface-name;
      }
    }
    voip {
      interface (all | [interface-name | access-ports]) {
        vlan vlan-name ;
        forwarding-class <assured-forwarding | best-effort | expedited-forwarding |
          network-control>;
      }
    }
  }
```

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure Ethernet switching options.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation

- Understanding Port Mirroring on J-EX Series Switches on page 3245
- Port Security for J-EX Series Switches Overview on page 2545
- Understanding BPDU Protection for STP, RSTP, and MSTP on J-EX Series Switches on page 1278
- Understanding Redundant Trunk Links on J-EX Series Switches on page 1049
- Understanding Storm Control on J-EX Series Switches on page 2511
- Understanding 802.1X and VoIP on J-EX Series Switches on page 2263
- Understanding Q-in-Q Tunneling on J-EX Series Switches on page 1051
- Understanding Unknown Unicast Forwarding on J-EX Series Switches on page 2512
- Understanding MAC Notification on J-EX Series Switches on page 1060

filter

Syntax filter (input | output) *filter-name*;

Hierarchy Level [edit vlans *vlan-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply a firewall filter to traffic coming into or exiting from the VLAN.

Default All incoming traffic is accepted unmodified to the VLAN, and all outgoing traffic is sent unmodified from the VLAN.

Options ***filter-name*** —Name of a firewall filter defined in a **filter** statement.

- **input**—Apply a firewall filter to VLAN ingress traffic.
- **output**—Apply a firewall filter to VLAN egress traffic.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Configuring Firewall Filters (CLI Procedure) on page 2779
- Configuring Firewall Filters (J-Web Procedure) on page 2784
- Firewall Filters for J-EX Series Switches Overview on page 2721

group-name

Syntax

```
group-name name {
  interface interface-name <primary>;
  interface interface-name;
}
```

Hierarchy Level [edit ethernet-switching-options redundant-trunk-group]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Create a redundant trunk group.

Options *name*—The name of the redundant trunk group. The group name must start with a letter and can consist of letters, numbers, dashes, and underscores.

The remaining options are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Redundant Trunk Links for Faster Recovery on page 1101
- Understanding Redundant Trunk Links on J-EX Series Switches on page 1049

gvrp

Syntax

```
gvrp {
  interface [interface-name] {
    disable;
  }
  join-timer milliseconds;
  leave-timer milliseconds;
  leaveall-timer milliseconds;
}
```

Hierarchy Level [edit protocols]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

NOTE: GVRP can be enabled only on trunk interfaces.

Description When GVRP is configured on a trunk interface, it ensures that the VLAN membership information on the trunk interface is updated as the switch's access interfaces become active or inactive in the configured VLANs.

The statements are explained separately.

Default GVRP is disabled by default.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show gvrp on page 1253**
- Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches on page 1070
- Example: Configure Automatic VLAN Administration Using GVRP on page 1087

instance-type

Syntax instance-type virtual-router

Hierarchy Level [edit routing-instances]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the type of routing instance.

Options **virtual-router**—A logical entity.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Using Virtual Routing Instances to Route Among VLANs on J-EX Series Switches on page 1112
- Configuring Virtual Routing Instances (CLI Procedure) on page 1142

interface

Syntax
```
interface (all | [interface-name]) {
    <enable | disable>;
}
```

Hierarchy Level [edit protocols gvrp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure GARP VLAN Registration Protocol (GVRP) for one or more interfaces.

Default By default, GVRP is disabled.

Options **all**—All interfaces.

interface-name—The list of interfaces to be configured for GVRP.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **show gvrp on page 1253**
- Example: Configure Automatic VLAN Administration Using GVRP on page 1087

interface (MVRP)

Syntax

```
interface (all | interface-name) {
    disable;
    join-timer milliseconds;
    leave-timer milliseconds;
    leaveall-timer milliseconds;
    registration (forbidden | normal);
}
```

Hierarchy Level [edit protocols mvrp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify interfaces on which to configure Multiple VLAN Registration Protocol (MVRP).

Default By default, MVRP is disabled.

Options **all**—All interfaces on the switch.

interface-name—Names of interface to be configured for MVRP.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Automatic VLAN Administration Using MVRP on J-EX Series Switches on page 1115
- Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure) on page 1147

interface

Syntax
interface *interface-name* <primary>;
interface *interface-name*;

Hierarchy Level
[edit **ethernet-switching-options redundant-trunk-group group-name** *name*]

Release Information
Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Configure a primary link and secondary link on trunk ports. If the primary link fails, the secondary link automatically takes over as the primary link without waiting for normal STP convergence.

Options
interface *interface-name*—A logical interface or an aggregated interface containing multiple ports.

primary—(Optional) Specify one of the interfaces in the redundant group as the primary link. The interface without this option is the secondary link in the redundant group. If a link is not specified as **primary**, the software compares the two links and selects the link with the highest port number as the active link. For example, if the two interfaces are **ge-0/1/0** and **ge-0/1/1**, the software assigns **ge-0/1/1** as the active link.

Required Privilege Level
routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation
- Example: Configuring Redundant Trunk Links for Faster Recovery on page 1101
- Understanding Redundant Trunk Links on J-EX Series Switches on page 1049

interface

Syntax
interface *interface-name*;

Hierarchy Level
[edit **routing-instances**]

Release Information
Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
For virtual routing instances, configure an interface.

Options
interface-name—Name of a Gigabit Ethernet interface.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Using Virtual Routing Instances to Route Among VLANs on J-EX Series Switches on page 1112
- Configuring Virtual Routing Instances (CLI Procedure) on page 1142
- Understanding Virtual Routing Instances on J-EX Series Switches on page 1048

interface

Syntax

```
interface interface-name {
  mapping (native (push | swap) | policy | tag (push | swap));
}
```

Hierarchy Level [edit vlans *vlan-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For a specific VLAN, configure an interface.

Options ***interface-name***—Name of a Gigabit Ethernet interface.

The remaining statement is explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch on page 1063
- Configuring VLANs for J-EX Series Switches (CLI Procedure) on page 1136
- Understanding Bridging and VLANs on J-EX Series Switches on page 1041
- Understanding Q-in-Q Tunneling on J-EX Series Switches on page 1051

interfaces

Syntax

```
interfaces interface-name {
  no-mac-learning;
}
```

Hierarchy Level [edit ethernet-switching-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure settings for interfaces that have been assigned to family **ethernet-switching**.

Options ***interface-name*** --Name of an interface that is configured for **family ethernet-switching**.

The remaining statement is explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation

- Understanding Q-in-Q Tunneling on J-EX Series Switches on page 1051

join-timer

Syntax join-timer *milliseconds*;

Hierarchy Level [edit protocols gvrp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For GARP VLAN Registration Protocol (GVRP), configure the maximum number of milliseconds interfaces must wait before sending VLAN advertisements.

Default 20 milliseconds

Options ***milliseconds*** —Number of milliseconds.
Default: 20 milliseconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **show gvrp on page 1253**
- Example: Configure Automatic VLAN Administration Using GVRP on page 1087

join-timer (MVRP)

Syntax join-timer *milliseconds*;

Hierarchy Level [edit protocols **mvrp interface** (all | *interface-name*)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the maximum number of milliseconds interfaces must wait before sending Multiple VLAN Registration Protocol (MVRP) protocol data units (PDUs).

Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.

Default 200 milliseconds

Options ***milliseconds***—Number of milliseconds that the interface must wait before sending MVRP PDUs.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **leave-timer on page 1204**
- **leaveall-timer on page 1206**
- Example: Configuring Automatic VLAN Administration Using MVRP on J-EX Series Switches on page 1115
- Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure) on page 1147

l3-interface

Syntax l3-interface vlan.*logical-interface-number*;

Hierarchy Level [edit **vlans** *vlan-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Associate a Layer 3 interface with the VLAN. Configure Layer 3 interfaces on trunk ports to allow the interface to transfer traffic between multiple VLANs. Within a VLAN, traffic is bridged, while across VLANs, traffic is routed.

Default No Layer 3 (routing) interface is associated with the VLAN.

Options vlan.*logical-interface-number*—Number of the logical interface defined with a **set interfaces vlan unit** command. For the logical interface number, use the same number you configure in the **unit** statement.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show ethernet-switching interfaces on page 997**
- **show vlans on page 1263**
- Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch on page 1063
- Example: Connecting an Access Switch to a Distribution Switch on page 1078
- Configuring Routed VLAN Interfaces (CLI Procedure) on page 1137
- Understanding Bridging and VLANs on J-EX Series Switches on page 1041

layer2-protocol-tunneling

Syntax

```
layer2-protocol-tunneling all | protocol-name {
    drop-threshold number;
    shutdown-threshold number;
}
```

Hierarchy Level [edit vlans *vlan-name* **dot1q-tunneling**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable Layer 2 protocol tunneling (L2PT) on the VLAN.

The remaining statements are explained separately.

Default L2PT is not enabled.

Options **all**—Enable all supported Layer 2 protocols.

protocol-name—Name of the Layer 2 protocol. Values are:

- **802.1x**—IEEE 802.1X authentication
- **802.3ah**—IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM)

  NOTE: If you enable L2PT for untagged OAM LFM packets, do not configure LFM on the corresponding access interface.

- **cdp**—Cisco Discovery Protocol
- **e-lmi**—Ethernet local management interface
- **gvrp**—GARP VLAN Registration Protocol
- **lacp**—Link Aggregation Control Protocol

  NOTE: If you enable L2PT for untagged LACP packets, do not configure LACP on the corresponding access interface.

- **llpd**—Link Layer Discovery Protocol
- **mmrp**—Multiple MAC Registration Protocol
- **mvrp**—Multiple VLAN Registration Protocol
- **stp**—Spanning Tree Protocol, Rapid Spanning Tree Protocol, and Multiple Spanning Tree Protocol
- **vstp**—VLAN Spanning Tree Protocol
- **vtp**—VLAN Trunking Protocol

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **show ethernet-switching layer2-protocol-tunneling interface on page 1234**
- **show ethernet-switching layer2-protocol-tunneling statistics on page 1236**
- **show ethernet-switching layer2-protocol-tunneling vlan on page 1239**
- Example: Configuring Layer 2 Protocol Tunneling on J-EX Series Switches on page 1126
- Configuring Layer 2 Protocol Tunneling on J-EX Series Switches (CLI Procedure) on page 1150

leave-timer

Syntax leave-timer *milliseconds*;

Hierarchy Level [edit protocols **gvrp**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For GARP VLAN Registration Protocol (GVRP), configure the number of milliseconds an interface waits after receiving a leave message before the interface leaves the VLAN specified in the message. If the interface receives a join message before the timer expires, the software keeps the interface in the VLAN.

Default 60 centiseconds

Options ***milliseconds*** —Number of milliseconds. At a minimum, the leave timer interval should be twice the join timer interval.
Default: 60 centiseconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **show gvrp on page 1253**
- Example: Configure Automatic VLAN Administration Using GVRP on page 1087

leave-timer (MVRP)

Syntax leave-timer *milliseconds*;

Hierarchy Level [edit protocols **mvrp interface** (all | *interface-name*)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For Multiple VLAN Registration Protocol (MVRP), configure the number of milliseconds the switch retains a VLAN in the Leave state before the VLAN is unregistered. If the interface receives a join message before this timer expires, the VLAN remains registered.

Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.

Default 1000 milliseconds

Options ***milliseconds***—Number of milliseconds that the switch retains a VLAN in the Leave state before the VLAN is unregistered. At a minimum, set the leave-timer interval at twice the join-timer interval.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **join-timer on page 1200**
- **leaveall-timer on page 1206**
- Example: Configuring Automatic VLAN Administration Using MVRP on J-EX Series Switches on page 1115
- Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure) on page 1147

leaveall-timer

Syntax leaveall-timer *milliseconds*;

Hierarchy Level [edit protocols **gvrp**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For GARP VLAN Registration Protocol (GVRP), configure the interval at which Leave All messages are sent on the interfaces. Leave All messages maintain current GVRP VLAN membership information in the network. A Leave All message instructs the port to change the GVRP state for all its VLANs to a leaving state and remove them unless a Join message is received before the leave timer expires.

Default 1000 centiseconds

Options ***milliseconds*** —Number of milliseconds.

Range: 5 times **leave-timer** value

Default: 1000 centiseconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show gvrp on page 1253**
- Example: Configure Automatic VLAN Administration Using GVRP on page 1087

leaveall-timer (MVRP)

Syntax leaveall-timer *milliseconds*;

Hierarchy Level [edit protocols **mvrp interface** (all | *interface-name*)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For Multiple VLAN Registration Protocol (MVRP), configure the interval at which the LeaveAll state operates on the interface.

Maintain default timer settings unless there is a compelling reason to change the settings. Modifying timers to inappropriate values might cause an imbalance in the operation of MVRP.

Default 10000 milliseconds

Options ***milliseconds***—Number of milliseconds between the sending of Leave All messages.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **join-timer on page 1200**
- **leave-timer on page 1204**
- Example: Configuring Automatic VLAN Administration Using MVRP on J-EX Series Switches on page 1115
- Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure) on page 1147

mac-limit

Syntax mac-limit *number*;

Hierarchy Level [edit **vlans** *vlan-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the number of MAC addresses allowed on a VLAN.

The number of MAC addresses allowed per VLAN varies between J-EX switches. The output of the **set vlans *vlan-name* mac-limit ?** provides the number of MAC addresses allowed on your J-EX switch.

Default MAC limit is disabled.

Options *number*—Maximum number of MAC addresses.
Range: 1 through 32768.

NOTE: Do not set **mac-limit** to 1. The first learned MAC address is often inserted into the forwarding database automatically (for instance, for routed VLAN Interfaces (RVIs), the first MAC address inserted into the forwarding database is the MAC address of the RVI. For aggregated Ethernet bundles using LACP, the first MAC address inserted into the forwarding database in the forwarding table is the source address of the protocol packet). The switch will therefore not learn MAC addresses other than the automatic addresses when the **mac-limit** is set to 1, and this will cause problems with MAC learning and forwarding.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show vlans on page 1263**
- Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch on page 1063
- Configuring MAC Table Aging (CLI Procedure) on page 1138
- Understanding Bridging and VLANs on J-EX Series Switches on page 1041

mac-notification

Syntax

```
mac-notification {
  notification-interval seconds;
}
```

Hierarchy Level [edit **ethernet-switching-options**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable MAC notification for a switch. If you configure this statement without setting a notification interval, MAC notification is enabled with the default MAC notification interval of 30 seconds.

The remaining statement is explained separately.

Default MAC notification is disabled by default.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring MAC Notification (CLI Procedure) on page 1151

mac-table-aging-time

Syntax mac-table-aging-time *seconds*;

Hierarchy Level [edit **ethernet-switching-options**],
[edit **vlans** *vlan-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define how long entries remain in the Ethernet switching table before expiring:

- If you specify this statement at the **[edit ethernet-switching-options]** hierarchy level, it applies to all VLANs on the switch.
- If you specify this statement at the **[edit vlans]** hierarchy level, it applies to the specified VLAN.

Default Entries remain in the Ethernet switching table for 300 seconds

Options ***seconds***—Time that entries remain in the Ethernet switching table before being removed.

- **Range**—60 through 1,000,000 seconds
- **Default**—300 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show ethernet-switching statistics aging on page 1244**
- Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch on page 1063
- Configuring MAC Table Aging (CLI Procedure) on page 1138
- Configuring VLANs for J-EX Series Switches (CLI Procedure) on page 1136
- Understanding Bridging and VLANs on J-EX Series Switches on page 1041

mapping

Syntax **mapping** (native (push | swap) | policy | *tag* (push | swap));

Hierarchy Level [edit **vlans** *vlan-name* **interface** *interface-name* ingress]:
[edit **vlans** *vlan-name* **interface** *interface-name* egress]:
[edit **vlans** *vlan-name* **interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Map a specific C-VLAN to an S-VLAN. By default, the received incoming or outgoing tag is replaced with the new tag.

This statement is also required if you are configuring firewall filters to map traffic from an interface to a VLAN. If you are configuring firewall filters to map traffic from an interface to a VLAN, the **mapping policy** option must be configured using this command. The firewall filter also has to be configured using the **vlan** action for a match condition in the firewall filter stanza for firewall filters to map traffic from an interface for a VLAN.

Options

native—Maps untagged and priority-tagged packets to an S-VLAN.

policy—Maps the interface to a firewall filter policy to an S-VLAN.

push—Retains the incoming tag and add an additional VLAN tag instead of replacing the original tag.

swap—Swaps the incoming VLAN tag with the VLAN ID tag of the S-VLAN. Use of this option is also referred to as VLAN ID translation.

tag—Retains the incoming 802.1Q tag on the interface.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring VLANs for J-EX Series Switches (CLI Procedure) on page 1136
- Understanding Q-in-Q Tunneling on J-EX Series Switches on page 1051
- Understanding Bridging and VLANs on J-EX Series Switches on page 1041

members

Syntax members [(all | *names* | *vlan-ids*)];

Hierarchy Level [edit **interfaces** *interface-name* **unit** *logical-unit-number* **family ethernet-switching vlan**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For trunk interfaces, configure the VLANs for which the interface can carry traffic.

TIP: To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type ? after vlan or vlans in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.

Options **all**—Specifies that this trunk interface is a member of all the VLANs that are configured on this switch. When a new VLAN is configured on the switch, this trunk interface automatically becomes a member of the VLAN.

NOTE: Each VLAN that is configured must have a specified VLAN ID when you attempt to commit the configuration; otherwise, the configuration commit fails. Also, all cannot be the name of a VLAN on the switch.

names—Name of one or more VLANs.

vlan-ids—Numeric identifier of one or more VLANs. For a series of tagged VLANs, specify a range; for example, **10-20** or **10-20 23 27-30**.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- **show ethernet-switching interfaces on page 997**
- **show vlans on page 1263**
- Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch on page 1063
- Example: Connecting an Access Switch to a Distribution Switch on page 1078
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 919
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 909
- Creating a Series of Tagged VLANs (CLI Procedure) on page 1140
- Understanding Bridging and VLANs on J-EX Series Switches on page 1041
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

mvrp

Syntax

```
mvrp {
  disable
  interface (all | interface-name) {
    disable;
    join-timer milliseconds;
    leave-timer milliseconds;
    leaveall-timer milliseconds;
    registration (forbidden | normal);
  }
  no-dynamic-vlan;
  traceoptions {
    file filename <files number > <size size> <no-stamp | world-readable |
      no-world-readable>;
    flag flag;
  }
}
```

Hierarchy Level [edit protocols]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure Multiple VLAN Registration Protocol (MVRP) on a trunk interface to ensure that the VLAN membership information on the trunk interface is updated as the switch's access interfaces become active or inactive in the configured VLANs.

The remaining statements are explained separately.

Default MVRP is disabled by default.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Automatic VLAN Administration Using MVRP on J-EX Series Switches on page 1115
- Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure) on page 1147

native-vlan-id

Syntax native-vlan-id *vlan-id*;

Hierarchy Level [edit **interfaces** *interface-name* unit 0 **family** ethernet-switching]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the VLAN identifier to associate with untagged packets received on the interface.

Options ***vlan-id***—Numeric identifier of the VLAN.
Range: 0 through 4095

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show vlans on page 1263**
- **show ethernet-switching interfaces on page 997**
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 919
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 909
- Understanding Bridging and VLANs on J-EX Series Switches on page 1041
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

no-dynamic-vlan

Syntax no-dynamic-vlan;

Hierarchy Level [edit protocols mvrp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable the dynamic creation of VLANs using Multiple VLAN Registration Protocol (MVRP) for interfaces participating in MVRP.

Dynamic VLAN configuration can be enabled on an interface independent of MVRP. The MVRP dynamic VLAN configuration setting does not override the interface configuration dynamic VLAN configuration setting. If dynamic VLAN creation is disabled on the interface in the interface configuration, no dynamic VLANs are created on the interface, including dynamic VLANs created using MVRP.

This option can only be applied globally; it cannot be applied per interface.

Default If MVRP is enabled, the dynamic creation of VLANs as a result of MVRP protocol exchange messages is enabled.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure) on page 1147

no-local-switching

Syntax no-local-switching

Hierarchy Level [edit vlans *vlan-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify that access ports in this VLAN domain do not forward packets to each other. You use this statement with primary VLANs and isolated secondary VLANs.

Required Privilege Level routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation
- Example: Configuring a Private VLAN on a J-EX Series Switch on page 1107
- Creating a Private VLAN (CLI Procedure) on page 1143

no-mac-learning

Syntax no-mac-learning;

Hierarchy Level [edit vlans *vlan-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disables MAC address learning for the specified VLAN.

Options There are no options to this statement.

Required Privilege Level routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation
- Configuring Q-in-Q Tunneling (CLI Procedure) on page 1144
- Understanding Q-in-Q Tunneling on J-EX Series Switches on page 1051

no-mac-learning

Syntax no-mac-learning;

Hierarchy Level [edit ethernet-switching-options interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable MAC address learning for the specified interface. Disabling MAC address learning on an interface disables learning for all the VLANs of which that interface is a member.

Options There are no options to this statement.

Required Privilege Level routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation
- Understanding Q-in-Q Tunneling on J-EX Series Switches on page 1051

notification-interval

Syntax notification-interval *seconds*;

Hierarchy Level [edit **ethernet-switching-options mac-notification**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the MAC notification interval for a switch.

The MAC notification interval is the amount of time the switch waits before sending learned or unlearned MAC address SNMP notifications to the network management server. For instance, if the MAC notification interval is set to 10, all of the MAC address addition and removal SNMP notifications will be sent to the network management system every 10 seconds.

Options *seconds*—The MAC notification interval, in seconds.

Range: 1 through 60

Default: 30

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring MAC Notification (CLI Procedure) on page 1151

port-mode

Syntax port-mode *mode*;

Hierarchy Level [edit **interfaces** *interface-name* **unit** *logical-unit-number* **family** ethernet-switching]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure whether an interface on the switch operates in access or trunk mode.

Default All switch interfaces are in access mode.

Options

access—Have the interface operate in access mode. In this mode, the interface can be in a single VLAN only. Access interfaces typically connect to network devices such as PCs, printers, IP telephones, and IP cameras.

trunk—Have the interface operate in trunk mode. In this mode, the interface can be in multiple VLANs and can multiplex traffic between different VLANs. Trunk interfaces typically connect to other switches and to routers on the LAN.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Connecting an Access Switch to a Distribution Switch on page 1078
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 919
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 909
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

primary-vlan

Syntax primary-vlan *vlan-name*

Hierarchy Level [edit **vlans** *vlan-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the primary VLAN for this community VLAN. The primary VLAN must be tagged, and the community VLAN must be untagged.

TIP: To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type ? after vlan or vlans in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.

Required Privilege Level routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation

- Example: Configuring a Private VLAN on a J-EX Series Switch on page 1107
- Creating a Private VLAN (CLI Procedure) on page 1143

redundant-trunk-group

Syntax

```
redundant-trunk-group {
  group-name name {
    interface interface-name <primary>;
    interface interface-name;
  }
}
```

Hierarchy Level [edit **ethernet-switching-options**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a primary link and secondary link on trunk ports. If the primary link fails, the secondary link automatically takes over without waiting for normal STP convergence.

Options The statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Redundant Trunk Links for Faster Recovery on page 1101
- Understanding Redundant Trunk Links on J-EX Series Switches on page 1049

registration

Syntax registration (forbidden | normal);

Hierarchy Level [edit protocols **mvrp interface** (all | *interface-name*)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specifies the Multiple VLAN Registration Protocol (MVRP) registration mode for the interface if MVRP is enabled.

Default **normal**

Options **forbidden**—The interface or interfaces do not register and do not participate in MVRP.

normal—The interface or interfaces accept MVRP messages and participate in MVRP.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Multiple VLAN Registration Protocol (MVRP) (CLI Procedure) on page 1147

routing-instances

Syntax
```
routing-instances routing-instance-name {
    instance-type virtual-router;
    interface interface-name;
}
```

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a virtual routing entity.

Options ***routing-instance-name***—Name for this routing instance.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Using Virtual Routing Instances to Route Among VLANs on J-EX Series Switches on page 1112
- Configuring Virtual Routing Instances (CLI Procedure) on page 1142

shutdown-threshold

Syntax shutdown-threshold *number*;

Hierarchy Level [edit **vlans** *vlan-name* dot1q-tunneling layer2-protocol-tunneling all | *protocol-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the maximum number of Layer 2 PDUs of the specified protocol that can be received per second on the interfaces in a specified VLAN before the interface is disabled. Once an interface is disabled, you must explicitly reenable it using the **clear ethernet-switching layer2-protocol-tunneling error** command. Otherwise, the interface remains disabled.

The shutdown threshold value must be greater than or equal to the drop threshold value. If the shutdown threshold value is less than the drop threshold value, the drop threshold value has no effect.

You can specify a shutdown threshold value without specifying a drop threshold value.

Default No shutdown threshold is specified.

Options ***number***—Maximum number of Layer 2 PDUs of the specified protocol that can be received per second on the interfaces in a specified VLAN before the interface is disabled.

Range: 1 through 1000

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **drop-threshold on page 1187**
- Example: Configuring Layer 2 Protocol Tunneling on J-EX Series Switches on page 1126
- Configuring Layer 2 Protocol Tunneling on J-EX Series Switches (CLI Procedure) on page 1150

vlan

Syntax

```
vlan {
  members [ (all | names | vlan-ids) ];
}
```

Hierarchy Level [edit **interfaces** *interface-name* **unit** *logical-unit-number* **family** ethernet-switching]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Bind an 802.1Q VLAN tag ID to a logical interface.

The remaining statement is explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- **show ethernet-switching interfaces on page 997**
- Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches on page 1070
- Configuring Routed VLAN Interfaces (CLI Procedure) on page 1137
- Understanding Bridging and VLANs on J-EX Series Switches on page 1041
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

vlan-id

Syntax vlan-id *number*;

Hierarchy Level [edit **vlans** *vlan-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure an 802.1Q tag to apply to all traffic that originates on the VLAN.

Default If you use the default factory configuration, all traffic originating on the VLAN is untagged and has a VLAN identifier of **0**.

Options ***number*** —VLAN tag identifier.

Range: 0 through 4093.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches on page 1070
- Understanding Bridging and VLANs on J-EX Series Switches on page 1041

vlan-range

Syntax vlan-range *vlan-id-low-vlan-id-high*;

Hierarchy Level [edit **vlans** *vlan-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure multiple VLANs. Each VLAN is assigned a VLAN ID number from the range.

Default None.

Options ***vlan-id-low-vlan-id-high*** —Specify the first and last VLAN ID number for the group of VLANs.

Required Privilege Level routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation

- Configuring VLANs for J-EX Series Switches (CLI Procedure) on page 1136
- Configuring VLANs for J-EX Series Switches (J-Web Procedure) on page 1133
- Configuring Routed VLAN Interfaces (CLI Procedure) on page 1137
- Understanding Bridging and VLANs on J-EX Series Switches on page 1041

vlans

Syntax

```
vlans {
  vlan-name {
    description text-description;
    dot1q-tunneling {
      customer-vlans (id | range)
      layer2-protocol-tunneling all | protocol-name {
        drop-threshold number;
        shutdown-threshold number;
      }
    }
    filter input filter-name;
    filter output filter-name;
    interface interface-name {
      mapping (native (push | swap) | policy | tag (push | swap));
    }
    l3-interface vlan.logical-interface-number;
    mac-limit number;
    mac-table-aging-time seconds;
    no-local-switching;
    no-mac-learning;
    primary-vlan vlan-name;
    vlan-id number;
    vlan-range vlan-id-low-vlan-id-high;
  }
}
```

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure VLAN properties on J-EX Series switches. The following configuration guidelines apply:

- Only private VLAN (PVLAN) firewall filters can be used when the VLAN is enabled for Q-in-Q tunneling.
- An S-VLAN tag is added to the packet if the VLAN is dot1q-tunneled and the packet is arriving from an access interface.
- You cannot use a firewall filter to assign a routed VLAN interface (RVI) to a VLAN.
- VLAN assignments performed using a firewall filter override all other VLAN assignments.

Default If you use the default factory configuration, all switch interfaces become part of the VLAN **default**.

Options ***vlan-name***—Name of the VLAN. The name can contain letters, numbers, hyphens (-), and periods (.) and can be up to 255 characters long.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation

- Configuring VLANs for J-EX Series Switches (CLI Procedure) on page 1136
- Configuring VLANs for J-EX Series Switches (J-Web Procedure) on page 1133
- Configuring Q-in-Q Tunneling (CLI Procedure) on page 1144
- Creating a Series of Tagged VLANs (CLI Procedure) on page 1140
- Configuring Routed VLAN Interfaces (CLI Procedure) on page 1137
- Understanding Q-in-Q Tunneling on J-EX Series Switches on page 1051
- Understanding Bridging and VLANs on J-EX Series Switches on page 1041

CHAPTER 63

Operational Mode Commands for Bridging and VLANs

clear ethernet-switching layer2-protocol-tunneling error

Syntax

clear ethernet-switching layer2-protocol-tunneling error
<interface *interface-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Clear Layer 2 protocol tunneling (L2PT) errors on one or more interfaces. If an interface has been disabled because the amount of Layer 2 protocol traffic exceeded the **shutdown-threshold** or because the switch has detected an error in the network topology or configuration, use this command to reenable the interface.

Options

none—Clears L2PT errors on all interfaces.

interface *interface-name*—(Optional) Clear L2PT errors on the specified interface.

Required Privilege Level

view

Related Documentation

- Example: Configuring Layer 2 Protocol Tunneling on J-EX Series Switches on page 1126
- Configuring Layer 2 Protocol Tunneling on J-EX Series Switches (CLI Procedure) on page 1150

List of Sample Output

clear ethernet-switching layer2-protocol-tunneling error

```
user@switch> clear ethernet-switching layer2-protocol-tunneling error
```

clear ethernet-switching layer2-protocol-tunneling error interface ge-0/1/1.0

```
user@switch> clear ethernet-switching layer2-protocol-tunneling error interface ge-0/1/1.0
```

clear ethernet-switching layer2-protocol-tunneling statistics

Syntax

clear ethernet-switching layer2-protocol-tunneling statistics
<interface *interface-name*>
<vlan *vlan-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Clear Layer 2 protocol tunneling (L2PT) statistics on one or more interfaces or VLANs.

Options

none—Clear L2PT statistics on all interfaces and VLANs.

interface *interface-name*—(Optional) Clear L2PT statistics on the specified interface.

vlan *vlan-name*—(Optional) Clear L2PT statistics on the specified VLAN.

Required Privilege Level

view

Related Documentation

- **show ethernet-switching layer2-protocol-tunneling statistics on page 1236**
- Example: Configuring Layer 2 Protocol Tunneling on J-EX Series Switches on page 1126
- Configuring Layer 2 Protocol Tunneling on J-EX Series Switches (CLI Procedure) on page 1150

List of Sample Output

clear ethernet-switching layer2-protocol-tunneling statistics on page 1227
clear ethernet-switching layer2-protocol-tunneling error interface ge-0/1/1.0 on page 1227
clear ethernet-switching layer2-protocol-tunneling error vlan v2 on page 1227

clear ethernet-switching layer2-protocol-tunneling statistics

```
user@switch> clear ethernet-switching layer2-protocol-tunneling statistics
```

clear ethernet-switching layer2-protocol-tunneling error interface ge-0/1/1.0

```
user@switch> clear ethernet-switching layer2-protocol-tunneling statistics interface ge-0/1/1.0
```

clear ethernet-switching layer2-protocol-tunneling error vlan v2

```
user@switch> clear ethernet-switching layer2-protocol-tunneling statistics vlan v2
```

clear ethernet-switching table

Syntax

clear ethernet-switching table
<interface *interface-name*>
<mac *mac-address*>
<management-vlan>
<vlan *vlan-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Clear learned entries, which are media access control (MAC) addresses, in the Ethernet switching table (also called the forwarding database table).

Options

none—Clear learned entries in the Ethernet switching table.

interface *interface-name*—(Optional) Clear all learned MAC addresses for the specified interface from the Ethernet switching table.

mac *mac-address*—(Optional) Clear the specified learned MAC address from the Ethernet switching table.

management-vlan—(Optional) Clear all MAC addresses learned for the management VLAN from the Ethernet switching table. Note that you do not specify a VLAN name because only one management VLAN exists.

vlan *vlan-name*—(Optional) Clear all MAC addresses learned for the specified VLAN from the Ethernet switching table.

Required Privilege Level

view

Related Documentation

- **show ethernet-switching table on page 1249**

List of Sample Output

clear ethernet-switching table on page 1228

Output Fields

This command produces no output.

clear ethernet-switching table

```
user@host> clear ethernet-switching table
```

clear gvrp statistics

Syntax clear gvrp statistics

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Clear GARP VLAN Registration Protocol (GVRP) statistics.

Required Privilege Level clear

Related Documentation
- **show spanning-tree statistics on page 1416**
- Example: Configure Automatic VLAN Administration Using GVRP on page 1087

List of Sample Output **clear gvrp statistics on page 1229**

clear gvrp statistics

```
user@switch> clear gvrp statistics
```

clear mvrp statistics

Syntax clear mvrp statistics <interface *interface-name*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Clear Multiple VLAN Registration Protocol (MVRP) statistics.

Options none—Clear all MVRP statistics.

interface *interface-name*—Clear the MVRP statistics on the specified interface.

Required Privilege Level clear

Related Documentation

- **show mvrp statistics on page 1260**
- Example: Configuring Automatic VLAN Administration Using MVRP on J-EX Series Switches on page 1115

List of Sample Output

clear mvrp statistics on page 1230
clear mvrp statistics interface ge-0/0/1.0 on page 1230

Output Fields When you enter this command, you are provided feedback on the status of your request.

clear mvrp statistics

```
user@switch> clear mvrp statistics
```

clear mvrp statistics interface ge-0/0/1.0

```
user@switch> clear mvrp statistics interface ge-0/0/1.0
```

show ethernet-switching interfaces

Syntax

show ethernet-switching interfaces
<brief | detail | summary>
<interface *interface-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

- **Blocking** field output updated.
- The default view updated to include information about 802.1Q-tags.
- The **detail** view updated to include information VLAN mapping.

Description

Display information about switched Ethernet interfaces.

Options

none—(Optional) Display brief information for Ethernet switching interfaces.

brief | detail | summary—(Optional) Display the specified level of output.

interface *interface-name*—(Optional) Display Ethernet switching information for a specific interface.

Required Privilege Level

view

Related Documentation

- **show ethernet-switching mac-learning-log on page 1241**
- **show ethernet-switching table on page 1249**
- Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure) on page 2516

List of Sample Output

show ethernet-switching interfaces on page 1232
show ethernet-switching interfaces ge-0/0/15 brief on page 1233
show ethernet-switching interfaces ge-0/0/2 detail (Blocked by RTG rtggroup) on page 1233
show ethernet-switching interfaces ge-0/0/15 detail (Blocked by STP) on page 1233
show ethernet-switching interfaces ge-0/0/17 detail (Disabled by bpdu-control) on page 1233
show ethernet-switching interfaces detail (C-VLAN to S-VLAN Mapping) on page 1233

Output Fields

Table 155 on page 1232 lists the output fields for the **show ethernet-switching interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 155: show ethernet-switching interfaces Output Fields

Field Name	Field Description	Level of Output
Interface	Name of a switching interface.	All levels
State	Interface state. Values are **up** and **down**.	none, **brief**, **detail**, **summary**
VLAN members	Name of a VLAN.	none, **brief**, **detail**, **summary**
Tag	Number of the 802.1Q-tag.	All levels
Tagging	Specifies whether the interface forwards 802.1Q-tagged or untagged traffic.	All levels
Blocking	The forwarding state of the interface: • **unblocked**—Traffic is forwarded on the interface. • **blocked**—Traffic is not being forwarded on the interface. • **Disabled by bpdu control**—The interface is disabled due to receiving BPDUs on a protected interface. If the **disable-timeout** statement has been included in the BPDU configuration, the interface automatically returns to service after the timer expires. • **blocked by RTG**—The specified redundant trunk group is disabled. • **blocked by STP**—The interface is disabled due to a spanning tree protocol error. • **MAC limit exceeded**—The interface is temporarily disabled due to a MAC limiting error. The disabled interface is automatically restored to service when the disable timeout expires. • **MAC move limit exceeded**—The interface is temporarily disabled due to a MAC move limiting error. The disabled interface is automatically restored to service when the disable timeout expires. • **Storm control in effect**—The interface is temporarily disabled due to a storm control error. The disabled interface is automatically restored to service when the disable timeout expires.	none, **brief**, **detail**, **summary**
Index	The VLAN index internal to Junos OS.	**detail**
mapping	The C-VLAN to S-VLAN mapping information: • **dot1q-tunneled**—The interface maps all traffic to the S-VLAN (all-in-one bundling). • **native**—The interface maps untagged and priority tagged packets to the S-VLAN. • **push**—The interface maps packets to a firewall filter to an S-VLAN. • **policy-mapped**—The interface maps packets to a specifically defined S-VLAN. • *integer*—The interface maps packets to the specified S-VLAN.	**detail**

show ethernet-switching interfaces

```
user@switch> show ethernet-switching interfaces

Interface    State  VLAN members        Tag   Tagging  Blocking
```

```
ae0.0          up      default                    untagged unblocked
ge-0/0/2.0     up      vlan300            300     untagged blocked by RTG (rtggroup)
ge-0/0/3.0     up      default                             blocked by STP
ge-0/0/4.0     down    default                             MAC limit exceeded
ge-0/0/5.0     down    default                             MAC move limit exceeded
ge-0/0/6.0     down    default                             Storm control in effect
ge-0/0/7.0     down    default                             unblocked
ge-0/0/13.0    up      default                    untagged unblocked
ge-0/0/14.0    up      vlan100             100    tagged   unblocked
                       vlan200             200    tagged   unblocked
ge-0/0/15.0    up      vlan100             100    tagged   blocked by STP
                       vlan200             200    tagged   blocked by STP
ge-0/0/16.0    down    default                    untagged unblocked
ge-0/0/17.0    down    vlan100             100    tagged   Disabled by bpdu-control

                       vlan200             200    tagged   Disabled by bpdu-control
```

show ethernet-switching interfaces ge-0/0/15 brief

```
user@switch> show ethernet-switching interfaces ge-0/0/15 brief
Interface   State  VLAN members   Tag   Tagging        Blocking

ge-0/0/15.0  up     vlan100        100   tagged         blocked by STP
                    vlan200        200   tagged         blocked by STP
```

show ethernet-switching interfaces ge-0/0/2 detail (Blocked by RTG rtggroup)

```
user@switch> show ethernet-switching interfaces ge-0/0/2 detail

Interface: ge-0/0/2.0, Index: 65, State: up, Port mode: Access
VLAN membership:
    vlan300, 802.1Q Tag: 300, untagged, msti-id: 0, blocked by RTG(rtggroup)
Number of MACs learned on IFL: 0
```

show ethernet-switching interfaces ge-0/0/15 detail (Blocked by STP)

```
user@switch> show ethernet-switching interfaces ge-0/0/15 detail

Interface: ge-0/0/15.0, Index: 70, State: up, Port mode: Trunk
VLAN membership:
    vlan100, 802.1Q Tag: 100, tagged, msti-id: 0, blocked by STP
    vlan200, 802.1Q Tag: 200, tagged, msti-id: 0, blocked by STP

Number of MACs learned on IFL: 0
```

show ethernet-switching interfaces ge-0/0/17 detail (Disabled by bpdu-control)

```
user@switch> show ethernet-switching interfaces ge-0/0/17 detail

Interface: ge-0/0/17.0, Index: 71, State: down, Port mode: Trunk
VLAN membership:
    vlan100, 802.1Q Tag: 100, tagged, msti-id: 1, Disabled by bpdu-control
    vlan200, 802.1Q Tag: 200, tagged, msti-id: 2, Disabled by bpdu-control
Number of MACs learned on IFL: 0
```

show ethernet-switching interfaces detail (C-VLAN to S-VLAN Mapping)

```
user@switch>show ethernet-switching interfaces ge-0/0/6.0 detail
Interface: ge-0/0/6.0, Index: 73, State: up, Port mode: Access
VLAN membership:
    map, 802.1Q Tag: 134, Mapped Tag: native, push, dot1q-tunneled, unblocked
    map, 802.1Q Tag: 134, Mapped Tag: 20, push, dot1q-tunneled, unblocked
```

show ethernet-switching layer2-protocol-tunneling interface

Syntax show ethernet-switching-layer2-protocol-tunneling interface <*interface-name*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display information about Layer 2 protocol tunneling (L2PT) on interfaces that have been configured for L2PT.

Options none—Display L2PT information about all interfaces on which L2PT is enabled.

interface-name—(Optional) Display L2PT information for the specified interface.

Required Privilege Level view

Related Documentation

- **show ethernet-switching layer2-protocol-tunneling statistics on page 1236**
- **show ethernet-switching layer2-protocol-tunneling vlan on page 1239**
- Configuring Layer 2 Protocol Tunneling on J-EX Series Switches (CLI Procedure) on page 1150

List of Sample Output **show ethernet-switching layer2-protocol-tunneling interface on page 1234**
show ethernet-switching layer2-protocol-tunneling interface ge-0/0/0.0 on page 1235

Output Fields Table 156 on page 1234 lists the output fields for the **show ethernet-switching layer2-protocol-tunneling interface** command. Output fields are listed in the approximate order in which they appear.

Table 156: show ethernet-switching layer2-protocol-tunneling interface Output Fields

Field Name	Field Description
Interface	Name of an interface on the switch.
Operation	Type of operation being performed on the interface. Values are **Encapsulation** and **Decapsulation**.
State	State of the interface. Values are **active** and **shutdown**.
Description	If the interface state is **shutdown**, displays why the interface is shut down. If the description says **Loop detected**, it means that the interface is an access interface that has received L2PT-enabled PDUs. Access interfaces should not receive L2PT-enabled PDUs. This scenario might mean that there is a loop in the network.

show ethernet-switching layer2-protocol-tunneling interface

```
user@switch> show ethernet-switching layer2-protocol-tunneling interface

Layer2 Protocol Tunneling information:
Interface     Operation       State      Description
ge-0/0/0.0    Encapsulation   Shutdown   Shutdown threshold exceeded
```

```
ge-0/0/1.0    Decapsulation   Shutdown   Loop detected
ge-0/0/2.0    Decapsulation   Active
```

show ethernet-switching layer2-protocol-tunneling interface ge-0/0/0.0

```
user@switch> show ethernet-switching layer2-protocol-tunneling interface ge-0/0/0.0

Layer2 Protocol Tunneling information:
Interface     Operation       State      Description
ge-0/0/0.0    Encapsulation   Shutdown   Shutdown threshold exceeded
```

show ethernet-switching layer2-protocol-tunneling statistics

Syntax show ethernet-switching-layer2-protocol-tunneling statistics
<interface *interface-name*>
<vlan *vlan-name*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display Layer 2 protocol tunneling (L2PT) statistics for Layer 2 PDU packets received by the switch.

NOTE: The show ethernet-switching-layer2-protocol-tunneling statistics command does not display L2PT statistics for Layer 2 PDU packets transmitted from the switch.

Options none—Display L2PT statistics for all interfaces on which you enabled L2PT.

<interface *interface-name*>—(Optional) Display L2PT statistics for the specified interface.

<vlan *vlan-name*>—(Optional) Display L2PT statistics for the specified VLAN.

Required Privilege Level view

Related Documentation

- **clear ethernet-switching layer2-protocol-tunneling statistics on page 1227**
- **show ethernet-switching layer2-protocol-tunneling interface on page 1234**
- **show ethernet-switching layer2-protocol-tunneling vlan on page 1239**
- **show vlans on page 1263**
- Example: Configuring Layer 2 Protocol Tunneling on J-EX Series Switches on page 1126
- Configuring Layer 2 Protocol Tunneling on J-EX Series Switches (CLI Procedure) on page 1150

List of Sample Output **show ethernet-switching layer2-protocol-tunneling statistics on page 1237**
show ethernet-switching layer2-protocol-tunneling statistics interface ge-0/0/0.0 on page 1237
show ethernet-switching layer2-protocol-tunneling statistics vlan v2 on page 1237

Output Fields Table 157 on page 1236 lists the output fields for the **show ethernet-switching layer2-protocol-tunneling statistics** command. Output fields are listed in the approximate order in which they appear.

Table 157: show ethernet-switching layer2-protocol-tunneling statistics Output Fields

VLAN	Field Description
VLAN	Name of a VLAN on which L2PT has been configured.

Table 157: show ethernet-switching layer2-protocol-tunneling statistics Output Fields *(continued)*

VLAN	Field Description
Interface	Name of an interface on which L2PT has been configured.
Protocol	Name of a protocol for which L2PT has been enabled. Values are **all**, **802.1x**, **802.3ah**, **cdp**, **e-lmi**, **gvrp**, **lacp**, **lldp**, **mmrp**, **mvrp**, **stp**, **vstp**, and **vtp**.
Operation	Type of operation being performed on the interface. Values are **Encapsulation** and **Decapsulation**.
Packets	Number of packets that have been encapsulated or decapsulated.
Drops	Number of packets that have exceeded the drop threshold and have been dropped.
Shutdowns	Number of times that packets have exceeded the shutdown threshold and the interface has been shut down.

show ethernet-switching layer2-protocol-tunneling statistics

```
user@switch> show ethernet-switching layer2-protocol-tunneling statistics

Layer2 Protocol Tunneling Statistics:
VLAN   Interface  Protocol  Operation     Packets Drops      Shutdowns
v1     ge-0/0/0.0 mvrp      Encapsulation 0       0          0
v1     ge-0/0/1.0 mvrp      Decapsulation 0       0          0
v1     ge-0/0/2.0 mvrp      Decapsulation 60634   0          0
v2     ge-0/0/0.0 cdp       Encapsulation 0       0          0
v2     ge-0/0/0.0 gvrp      Encapsulation 0       0          0
v2     ge-0/0/0.0 lldp      Encapsulation 0       0          0
```

show ethernet-switching layer2-protocol-tunneling statistics interface ge-0/0/0.0

```
user@switch> show ethernet-switching layer2-protocol-tunneling statistics interface ge-0/0/0.0

Layer2 Protocol Tunneling Statistics:
VLAN    Interface  Protocol Operation      Packets  Drops     Shutdowns
v1      ge-0/0/0.0 mvrp     Encapsulation  0        0          0
v2      ge-0/0/0.0 cdp      Encapsulation  0        0          0
v2      ge-0/0/0.0 gvrp     Encapsulation  0        0          0
v2      ge-0/0/0.0 lldp     Encapsulation  0        0          0
v2      ge-0/0/0.0 mvrp     Encapsulation  0        0          0
v2      ge-0/0/0.0 stp      Encapsulation  0        0          0
v2      ge-0/0/0.0 vtp      Encapsulation  0        0          0
v2      ge-0/0/0.0 vstp     Encapsulation  0        0          0
```

show ethernet-switching layer2-protocol-tunneling statistics vlan v2

```
user@switch> show ethernet-switching layer2-protocol-tunneling statistics vlan v2

Layer2 Protocol Tunneling Statistics:
VLAN    Interface   Protocol Operation       Packets  Drops   Shutdowns
v2      ge-0/0/0.0  cdp      Encapsulation   0        0       0
v2      ge-0/0/0.0  gvrp     Encapsulation   0        0       0
v2      ge-0/0/0.0  lldp     Encapsulation   0        0       0
v2      ge-0/0/0.0  mvrp     Encapsulation   0        0       0
v2      ge-0/0/0.0  stp      Encapsulation   0        0       0
v2      ge-0/0/0.0  vtp      Encapsulation   0        0       0
v2      ge-0/0/0.0  vstp     Encapsulation   0        0       0
v2      ge-0/0/1.0  cdp      Decapsulation   0        0       0
v2      ge-0/0/1.0  gvrp     Decapsulation   0        0       0
v2      ge-0/0/1.0  lldp     Decapsulation   0        0       0
```

```
v2      ge-0/0/1.0  mvrp     Decapsulation   0        0       0
v2      ge-0/0/1.0  stp      Decapsulation   0        0       0
v2      ge-0/0/1.0  vtp      Decapsulation   0        0       0
```

show ethernet-switching layer2-protocol-tunneling vlan

Syntax show ethernet-switching-layer2-protocol-tunneling vlan <*vlan-name*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display information about Layer 2 protocol tunneling (L2PT) on VLANs that have been configured for L2PT.

Options none—Display information about L2PT for the VLANs on which you have configured L2PT.

vlan-name—(Optional) Display information about L2PT for the specified VLAN.

Required Privilege Level view

Related Documentation

- **show ethernet-switching layer2-protocol-tunneling interface on page 1234**
- **show ethernet-switching layer2-protocol-tunneling statistics on page 1236**
- **show vlans on page 1263**
- Example: Configuring Layer 2 Protocol Tunneling on J-EX Series Switches on page 1126
- Configuring Layer 2 Protocol Tunneling on J-EX Series Switches (CLI Procedure) on page 1150

List of Sample Output **show ethernet-switching layer2-protocol-tunneling vlan on page 1239**
show ethernet-switching layer2-protocol-tunneling vlan v2 on page 1240

Output Fields Table 158 on page 1239 lists the output fields for the **show ethernet-switching layer2-protocol-tunneling vlan** command. Output fields are listed in the approximate order in which they appear.

Table 158: show ethernet-switching layer2-protocol-tunneling vlan Output Fields

Field Name	Field Description
VLAN	Name of the VLAN on which L2PT has been configured.
Protocol	Name of a protocol for which L2PT has been enabled. Values are **all**, **802.1x**, **802.3ah**, **cdp**, **e-lmi**, **gvrp**, **lacp**, **lldp**, **mmrp**, **mvrp**, **stp**, **vstp**, and **vtp**.
Drop Threshold	Maximum number of Layer 2 PDUs of the specified protocol that can be received per second on the VLAN before the switch begins dropping the Layer 2 PDUs.
Shutdown Threshold	Maximum number of Layer 2 PDUs of the specified protocol that can be received per second on the VLAN before the interface is disabled.

show ethernet-switching

```
user@switch> show ethernet-switching layer2-protocol-tunneling vlan

Layer2 Protocol Tunneling VLAN information:
```

layer2-protocol-tunneling vlan

```
VLAN            Protocol        Drop        Shutdown
                                Threshold   Threshold
v1              mvrp            100         200
v2              cdp             0           0
v2              cdp             0           0
v2              gvrp            0           0
```

show ethernet-switching layer2-protocol-tunneling vlan v2

```
user@switch> show ethernet-switching layer2-protocol-tunneling vlan v2

Layer2 Protocol Tunneling VLAN information:
VLAN            Protocol        Drop        Shutdown
                                Threshold   Threshold
v2              cdp             0           0
v2              cdp             0           0
v2              gvrp            0           0
```

show ethernet-switching mac-learning-log

Syntax show ethernet-switching mac-learning-log

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Displays the event log of learned MAC addresses.

Required Privilege Level view

Related Documentation

- **show ethernet-switching table on page 1249**
- **show ethernet-switching interfaces on page 997**
- Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch on page 1063
- Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches on page 1070
- Example: Configure Automatic VLAN Administration Using GVRP on page 1087
- Example: Connecting an Access Switch to a Distribution Switch on page 1078

List of Sample Output **show ethernet-switching mac-learning-log on page 1241**

Output Fields Table 159 on page 1241 lists the output fields for the **show ethernet-switching mac-learning-log** command. Output fields are listed in the approximate order in which they appear.

Table 159: show ethernet-switching mac-learning-log Output Fields

Field Name	Field Description
Date and Time	Timestamp when the MAC address was added or deleted from the log.
vlan_name	VLAN name. A value defined by the user for all user-configured VLANs.
MAC	Learned MAC address.
Deleted | Added	MAC address deleted or added to the MAC learning log.
Blocking	The forwarding state of the interface: • **blocked**—Traffic is not being forwarded on the interface. • **unblocked**—Traffic is forwarded on the interface.

show ethernet-switching mac-learning-log

```
user@switch> show ethernet-switching mac-learning-log
Mon Feb 25 08:07:05 2008
 vlan_name v1 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name v9 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name HR_vlan mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
```

```
 vlan_name v3 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name v12 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name v13 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name sales_vlan mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name employee1 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name employee2 mac 00:00:00:00:00:00 was deleted
Mon Feb 25 08:07:05 2008
 vlan_name v3 mac 00:00:00:00:00:00 was added
Mon Feb 25 08:07:05 2008
 vlan_name HR_vlan mac 00:00:00:00:00:00 was added
Mon Feb 25 08:07:05 2008
 vlan_name employee2 mac 00:00:00:00:00:00 was added
Mon Feb 25 08:07:05 2008
 vlan_name employee1 mac 00:00:00:00:00:00 was added
Mon Feb 25 08:07:05 2008
 vlan_name employee2 mac 00:00:05:00:00:05 was learned
Mon Feb 25 08:07:05 2008
 vlan_name employee1 mac 00:30:48:90:54:89 was learned
Mon Feb 25 08:07:05 2008
 vlan_name HR_vlan mac 00:00:5e:00:01:00 was learned
Mon Feb 25 08:07:05 2008
 vlan_name sales_vlan mac 00:00:5e:00:01:08 was learned
[output truncated]
```

show ethernet-switching mac-notification

Syntax show ethernet-switching mac-notification

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display information about MAC notification.

Required Privilege Level view

Related Documentation
- Verifying That MAC Notification Is Working Properly on page 1163

Output Fields Table 160 on page 1243 lists the output fields for the **show ethernet-switching mac-notification** command. Output fields are listed in the order in which they appear.

Table 160: show ethernet-switching mac-notification Output Fields

Field Name	Field Description
Notification Status	Displays the MAC notification status: • **Enabled**—MAC notification is enabled. • **Disabled**—MAC notification is disabled.
Notification Interval	Displays the MAC notification interval in seconds.

show ethernet-switching mac-notification (MAC Notification Enabled)

```
user@switch> show ethernet-switching mac-notification
Notification Status                : Enabled
Notification Interval              : 30
```

show ethernet-switching mac-notification (MAC Notification Disabled)

```
user@switch> show ethernet-switching mac-notification
Notification Status                : Disabled
Notification Interval              : 0
```

show ethernet-switching statistics aging

Syntax show ethernet-switching statistics aging

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display media access control (MAC) aging statistics.

Options none—(Optional) Display MAC aging statistics.

brief | detail—(Optional) Display the specified level of output.

Required Privilege Level view

Related Documentation

- **show ethernet-switching statistics mac-learning on page 1246**
- Configuring MAC Table Aging (CLI Procedure) on page 1138

List of Sample Output **show ethernet-switching statistics aging on page 1244**

Output Fields Table 161 on page 1244 lists the output fields for the **show ethernet-switching statistics aging** command. Output fields are listed in the approximate order in which they appear.

Table 161: show ethernet-switching statistics aging Output Fields

Field Name	Field Description	Level of Output
Total age messages received	Total number of aging messages received from the hardware.	All levels
Immediate aging	Aging message indicating that the entry should be removed immediately.	All levels
MAC address seen	Aging message indicating that the MAC address has been detected by hardware and that the aging timer should be stopped.	All levels
MAC address not seen	Aging message indicating that the MAC address has not been detected by the hardware and that the aging timer should be started.	All levels
Error age messages	The received aging message contains the following errors: • **Invalid VLAN**—The VLAN of the packet does not exist. • **No such entry**—The MAC address and VLAN pair provided by the aging message does not exist. • **Static entry**—An unsuccessful attempt was made to age out a static MAC entry.	All levels

show ethernet-switching statistics aging

```
user@switch> show ethernet-switching statistics aging

Total age messages received: 0
  Immediate aging: 0, MAC address seen: 0, MAC address not seen: 0
```

```
Error age messages: 0
  Invalid VLAN: 0, No such entry: 0, Static entry: 0
```

show ethernet-switching statistics mac-learning

Syntax show ethernet-switching statistics mac-learning

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display media access control (MAC) learning statistics.

Options none—(Optional) Display MAC learning statistics for all interfaces.

brief | detail—(Optional) Display the specified level of output.

interface *interface-name* —(Optional) Display MAC learning statistics for the specified interface.

Required Privilege Level view

Related Documentation

- **show ethernet-switching statistics aging on page 1244**
- **show ethernet-switching mac-learning-log on page 1241**
- **show ethernet-switching table on page 1249**
- **show ethernet-switching interfaces on page 997**
- Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch on page 1063
- Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches on page 1070
- Example: Configure Automatic VLAN Administration Using GVRP on page 1087

List of Sample Output

show ethernet-switching statistics mac-learning on page 1247
show ethernet-switching statistics mac-learning detail on page 1247
show ethernet-switching statistics mac-learning interface ge-0/0/1 on page 1248

Output Fields Table 162 on page 1246 lists the output fields for the **show ethernet-switching statistics mac-learning** command. Output fields are listed in the approximate order in which they appear.

Table 162: show ethernet-switching statistics mac-learning Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface for which statistics are being reported.	All levels
Learning message from local packets	MAC learning message generated due to packets coming in on the management interface.	All levels
Learning message from transit packets	MAC learning message generated due to packets coming in on network interfaces.	All levels

Table 162: show ethernet-switching statistics mac-learning Output Fields *(continued)*

Field Name	Field Description	Level of Output
Learning message with error	MAC learning messages received with errors: • **Invalid VLAN**—The VLAN of the packet does not exist. • **Invalid MAC**—The MAC address is either NULL or a multicast MAC address. • **Security violation**—The MAC address is not an allowed MAC address. • **Interface down**—The MAC address is learned on an interface that is down. • **Incorrect membership**—The MAC address is learned on an interface that is not a member of the VLAN. • **Interface limit**—The number of MAC addresses learned on the interface has exceeded the limit. • **MAC move limit**—This MAC address has moved among multiple interfaces too many times in a given interval. • **VLAN limit**—The number of MAC addresses learned on the VLAN has exceeded the limit. • **Invalid VLAN index**—The VLAN of the packet, while configured, does not yet exist in the kernel. • **Interface not learning**—The MAC address is learned on an interface that does not yet allow learning—for example, the interface is blocked. • **No nexthop**—The MAC address is learned on an interface that does not have a unicast next hop. • **MAC learning disabled**—The MAC address is learned on an interface on which MAC learning has been disabled. • **Others**—The message contains some other error.	All levels

show ethernet-switching statistics mac-learning

```
user@switch> show ethernet-switching statistics mac-learning

Learning stats: 0 learn msg rcvd, 0 error
  Interface         Local pkts        Transit pkts      Error
  ge-0/0/0.0        0                 0                 0
  ge-0/0/1.0        0                 0                 0
  ge-0/0/2.0        0                 0                 0
  ge-0/0/3.0        0                 0                 0
```

show ethernet-switching statistics mac-learning detail

```
user@switch> show ethernet-switching statistics mac-learning detail
Learning stats: 0 learn msg rcvd, 0 error

    Interface: ge-0/0/0.0
    Learning message from local packets:   0
    Learning message from transit packets: 1
    Learning message with error:           0
      Invalid VLAN:               0       Invalid MAC:                0
      Security violation:         0       Interface down:             0
      Incorrect membership:       0       Interface limit:            0
      MAC move limit:             0       VLAN limit:                 0
      Invalid VLAN index:         0       Interface not learning:     0
      No nexthop:                 0       MAC learning disabled:      0
      Others:                     0

    Interface: ge-0/0/1.0
    Learning message from local packets:   0
    Learning message from transit packets: 2
```

```
  Learning message with error:           0
    Invalid VLAN:                0       Invalid MAC:                   0
    Security violation:          0       Interface down:                0
    Incorrect membership:        0       Interface limit:               0
    MAC move limit:              0       VLAN limit:                    0
    Invalid VLAN index:          0       Interface not learning:        0
    No nexthop:                  0       MAC learning disabled:         0
    Others:                      0
```

show ethernet-switching statistics mac-learning interface ge-0/0/1

```
user@switch> show ethernet-switching statistics mac-learning interface ge-0/0/1
Interface         Local pkts        Transit pkts      Error
ge-0/0/1.0        0                 1                 1
```

show ethernet-switching table

Syntax
show ethernet-switching table
<brief | detail | extensive | summary>
<interface *interface-name*>
<management-vlan>
<sort-by (*name* | *tag*)
<vlan (*vlan-name*)>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Displays the Ethernet switching table.

Options
none—(Optional) Display brief information about the Ethernet switching table.

brief | detail | extensive | summary—(Optional) Display the specified level of output.

management-vlan—(Optional) Display the Ethernet switching table for a management VLAN.

interface-name—(Optional) Display the Ethernet switching table for a specific interface.

sort-by (*name* | *tag*)—(Optional) Display VLANs in ascending order of VLAN IDs or VLAN names.

vlan *vlan-name*—(Optional) Display the Ethernet switching table for a specific VLAN.

Required Privilege Level
view

Related Documentation
- Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch on page 1063
- Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches on page 1070
- Example: Configure Automatic VLAN Administration Using GVRP on page 1087
- Example: Setting Up Q-in-Q Tunneling on J-EX Series Switches on page 1105

List of Sample Output
show ethernet-switching table on page 1250
show ethernet-switching table brief on page 1251
show ethernet-switching table detail on page 1251
show ethernet-switching table extensive on page 1252
show ethernet-switching table interface ge-0/0/1 on page 1252

Output Fields
Table 163 on page 1249 lists the output fields for the **show ethernet-switching table** command. Output fields are listed in the approximate order in which they appear.

Table 163: show ethernet-switching table Output Fields

Field Name	Field Description	Level of Output
VLAN	The name of a VLAN.	All levels

Table 163: show ethernet-switching table Output Fields *(continued)*

Field Name	Field Description	Level of Output
Tag	The VLAN ID tag name or number.	**extensive**
MAC or **MAC address**	The MAC address associated with the VLAN.	All levels
Type	The type of MAC address. Values are: • **static**—The MAC address is manually created. • **learn**—The MAC address is learned dynamically from a packet's source MAC address. • **flood**—The MAC address is unknown and flooded to all members.	All levels
Age	The time remaining before the entry ages out and is removed from the Ethernet switching table.	All levels
Interfaces	Interface associated with learned MAC addresses or **All-members** (flood entry).	All levels
Learned	For learned entries, the time which the entry was added to the Ethernet-switching table.	**detail**, **extensive**
Nexthop index	The nexthop index number.	**detail**, **extensive**

show ethernet-switching table

```
user@switch> show ethernet-switching table
Ethernet-switching table: 57 entries, 17 learned
  VLAN              MAC address       Type         Age Interfaces
  F2                *                 Flood          - All-members
  F2                00:00:05:00:00:03 Learn          0 ge-0/0/44.0
  F2                00:19:e2:50:7d:e0 Static         - Router
  Linux             *                 Flood          - All-members
  Linux             00:19:e2:50:7d:e0 Static         - Router
  Linux             00:30:48:90:54:89 Learn          0 ge-0/0/47.0
  T1                *                 Flood          - All-members
  T1                00:00:05:00:00:01 Learn          0 ge-0/0/46.0
  T1                00:00:5e:00:01:00 Static         - Router
  T1                00:19:e2:50:63:e0 Learn          0 ge-0/0/46.0
  T1                00:19:e2:50:7d:e0 Static         - Router
  T10               *                 Flood          - All-members
  T10               00:00:5e:00:01:09 Static         - Router
  T10               00:19:e2:50:63:e0 Learn          0 ge-0/0/46.0
  T10               00:19:e2:50:7d:e0 Static         - Router
  T111              *                 Flood          - All-members
  T111              00:19:e2:50:63:e0 Learn          0 ge-0/0/15.0
  T111              00:19:e2:50:7d:e0 Static         - Router
  T111              00:19:e2:50:ac:00 Learn          0 ge-0/0/15.0
  T2                *                 Flood          - All-members
  T2                00:00:5e:00:01:01 Static         - Router
  T2                00:19:e2:50:63:e0 Learn          0 ge-0/0/46.0
  T2                00:19:e2:50:7d:e0 Static         - Router
  T3                *                 Flood          - All-members
  T3                00:00:5e:00:01:02 Static         - Router
  T3                00:19:e2:50:63:e0 Learn          0 ge-0/0/46.0
  T3                00:19:e2:50:7d:e0 Static         - Router
  T4                *                 Flood          - All-members
```

```
  T4                00:00:5e:00:01:03 Static         - Router
  T4                00:19:e2:50:63:e0 Learn          0 ge-0/0/46.0
[output truncated]
```

show ethernet-switching table brief

```
user@switch> show ethernet-switching table brief
Ethernet-switching table: 57 entries, 17 learned
  VLAN              MAC address       Type          Age Interfaces
  F2                *                 Flood           - All-members
  F2                00:00:05:00:00:03 Learn           0 ge-0/0/44.0
  F2                00:19:e2:50:7d:e0 Static          - Router
  Linux             *                 Flood           - All-members
  Linux             00:19:e2:50:7d:e0 Static          - Router
  Linux             00:30:48:90:54:89 Learn           0 ge-0/0/47.0
  T1                *                 Flood           - All-members
  T1                00:00:05:00:00:01 Learn           0 ge-0/0/46.0
  T1                00:00:5e:00:01:00 Static          - Router
  T1                00:19:e2:50:63:e0 Learn           0 ge-0/0/46.0
  T1                00:19:e2:50:7d:e0 Static          - Router
  T10               *                 Flood           - All-members
  T10               00:00:5e:00:01:09 Static          - Router
  T10               00:19:e2:50:63:e0 Learn           0 ge-0/0/46.0
  T10               00:19:e2:50:7d:e0 Static          - Router
  T111              *                 Flood           - All-members
  T111              00:19:e2:50:63:e0 Learn           0 ge-0/0/15.0
  T111              00:19:e2:50:7d:e0 Static          - Router
  T111              00:19:e2:50:ac:00 Learn           0 ge-0/0/15.0
  T2                *                 Flood           - All-members
  T2                00:00:5e:00:01:01 Static          - Router
  T2                00:19:e2:50:63:e0 Learn           0 ge-0/0/46.0
  T2                00:19:e2:50:7d:e0 Static          - Router
  T3                *                 Flood           - All-members
  T3                00:00:5e:00:01:02 Static          - Router
  T3                00:19:e2:50:63:e0 Learn           0 ge-0/0/46.0
  T3                00:19:e2:50:7d:e0 Static          - Router
  T4                *                 Flood           - All-members
  T4                00:00:5e:00:01:03 Static          - Router
  T4                00:19:e2:50:63:e0 Learn           0 ge-0/0/46.0
[output truncated]
```

show ethernet-switching table detail

```
user@switch> show ethernet-switching table detail
Ethernet-switching table: 5 entries, 2 learned
 VLAN: default, Tag: 0, MAC: *, Interface: All-members
  Interfaces:
   ge-0/0/11.0, ge-0/0/20.0, ge-0/0/30.0, ge-0/0/36.0, ge-0/0/3.0
  Type: Flood
  Nexthop index: 1307

 VLAN: default, Tag: 0, MAC: 00:1f:12:30:b8:83, Interface: ge-0/0/3.0
   Type: Learn, Age: 0, Learned: 20:09:26
   Nexthop index: 1315

 VLAN: v1, Tag: 101, MAC: *, Interface: All-members
  Interfaces:
   ge-0/0/31.0
  Type: Flood
  Nexthop index: 1313

 VLAN: v1, Tag: 101, MAC: 00:1f:12:30:b8:89, Interface: ge-0/0/31.0
  Type: Learn, Age: 0, Learned: 20:09:25
  Nexthop index: 1312
```

```
 VLAN: v2, Tag: 102, MAC: *, Interface: All-members
  Interfaces:
   ae0.0
  Type: Flood
  Nexthop index: 1317
```

show ethernet-switching table extensive

```
user@switch> show ethernet-switching table extensive
Ethernet-switching table: 3 entries, 1 learned

    VLAN: v1, Tag: 10, MAC: *, Interface: All-members
    Interfaces:
                ge-0/0/14.0, ge-0/0/1.0, ge-0/0/2.0, ge-0/0/3.0, ge-0/0/4.0,
                ge-0/0/5.0, ge-0/0/6.0, ge-0/0/7.0, ge-0/0/8.0, ge-0/0/10.0,
                ge-0/0/0.0
    Type: Flood
    Nexthop index: 567

    VLAN: v1, Tag: 10, MAC: 00:21:59:c6:93:22, Interface: Router
    Type: Static
    Nexthop index: 0

    VLAN: v1, Tag: 10, MAC: 00:21:59:c9:9a:4e, Interface: ge-0/0/14.0
    Type: Learn, Age: 0, Learned: 18:40:50
    Nexthop index: 564
```

show ethernet-switching table interface ge-0/0/1

```
user@switch> show ethernet-switching table interface ge-0/0/1
Ethernet-switching table: 1 unicast entries
  VLAN              MAC address       Type         Age Interfaces
  V1                *                 Flood          - All-members
  V1                00:00:05:00:00:05 Learn          0 ge-0/0/1.0
```

show gvrp

Syntax show gvrp

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display GARP VLAN Registration Protocol (GVRP) information.

Options none—Displays all GVRP configuration attributes.

interface *interface-name* —(Optional) Displays GVRP statistics for a specific interface only.

Required Privilege Level view

Related Documentation

- **show gvrp statistics on page 1255**
- Example: Configure Automatic VLAN Administration Using GVRP on page 1087

List of Sample Output **show gvrp on page 1253**

Output Fields Table 164 on page 1253 lists the output fields for the **show gvrp** command. Output fields are listed in the approximate order in which they appear.

Table 164: show gvrp Output Fields

Field Name	Field Description
Global GVRP Configuration	Displays global GVRP information: • **GVRP status**—Displays whether GVRP is **enabled**or **disabled**. • **Join**—The maximum number of milliseconds the interfaces must wait before sending VLAN advertisements. • **Leave**— The number of milliseconds an interface must wait after receiving a Leave message to remove the interface from the VLAN specified in the message. • **Leaveall**—The interval at which Leave All messages are sent on interfaces. Leave all messages maintain current GVRP VLAN membership information in the network.
Interface based configuration	Displays interface-specific GVRP information: • **Interface**—The interface on which GVRP is configured.. • **GVRP status**—Displays whether GVRP is **enabled** or **disabled**.

show gvrp

```
user@switch> show gvrp

Global GVRP configuration
  GVRP status   : Enabled
  GVRP timers (ms)
    Join        : 40
    Leave       : 120
    Leaveall    : 2000
```

```
Interface based configuration:
Interface  GVRP status
---------- -----------
ge-0/0/0.0 Enabled
```

show gvrp statistics

Syntax show gvrp statistics

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display Generic VLAN Registration Protocol (GVRP) statistics in the form of GARP Information Propagation (GIP) messages.

Required Privilege Level clear

Related Documentation
- **show gvrp on page 1253**
- Example: Configure Automatic VLAN Administration Using GVRP on page 1087

List of Sample Output **show gvrp statistics on page 1256**

Output Fields Table 165 on page 1255 lists the output fields for the **show gvrp statistics** command. Output fields are listed in the approximate order in which they appear.

Table 165: show gvrp statistics Output Fields

Field Name	Field Description
Join Empty received	Number of GIP **Join Empty** messages received on the switch.
Join In received	Number of GIP **Join In** messages received on the switch.
Empty received	Number of GIP **Empty** messages received on the switch.
Leave In received	Number of GIP **Leave In** messages received on the switch.
Leave Empty received	Number of GIP **Leave Empty** messages received on the switch.
Leave All received	Number of GIP **Leave All** messages received on the switch.
Join Empty transmitted	Number of GIP **Join Empty** messages sent from the switch.
Join In transmitted	Number of GIP **Join In** messages sent from the switch.
Empty transmitted	Number of GIP **Empty** messages sent from the switch.
Leave In transmitted	Number of GIP **Leave In** messages sent from the switch.
Leave Empty transmitted	Number of GIP **Leave Empty** messages sent from the switch.
Leave All transmitted	Number of GIP **Leave All** messages sent from the switch.

show gvrp statistics

```
user@switch> show gvrp statistics
GVRP statistics
  Join Empty received           : 0
  Join In received              : 12
  Empty received                : 0
  Leave In received             : 0
  Leave Empty received          : 0
  Leave All received            : 0
  Join Empty transmitted        : 0
  Join In transmitted           : 48
  Empty transmitted             : 4
  Leave In transmitted          : 0
  Leave Empty transmitted       : 0
  Leave All transmitted         : 4
```

show mvrp

Syntax show mvrp

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display Multiple VLAN Registration Protocol (MVRP) configuration information.

Required Privilege Level view

Related Documentation

- **show mvrp statistics on page 1260**
- Example: Configuring Automatic VLAN Administration Using MVRP on J-EX Series Switches on page 1115
- Verifying That MVRP Is Working Correctly on page 1162

List of Sample Output **show mvrp on page 1257**

Output Fields Table 166 on page 1257 lists the output fields for the **show mvrp** command. Output fields are listed in the approximate order in which they appear.

Table 166: show mvrp Output Fields

Field Name	Field Description
Global MVRP configuration	Displays global MVRP information: • **MVRP status**—Displays whether MVRP is **Enabled** or **Disabled**. • **MVRP dynamic vlan creation**—Displays whether global MVRP dynamic VLAN creation is **Dnabled** or **Disabled**.
MVRP Timers (ms)	Displays MVRP timer information: • **Interface**—The interface on which MVRP is configured. • **Join**—The maximum number of milliseconds the interfaces must wait before sending VLAN advertisements. • **Leave**—The number of milliseconds an interface must wait after receiving a Leave message to remove the interface from the VLAN specified in the message. • **LeaveAll**—The interval at which LeaveAll messages are sent on interfaces. LeaveAll messages maintain current MVRP VLAN membership information in the network.
Interface based configuration	Displays interface-specific MVRP information: • **Interface**—The interface on which MVRP is configured. • **Status**—Displays whether MVRP is **Enabled** or **Disabled**. • **Registration**—Displays whether registration for the interface is **Forbidden** or **Normal**. • **Dynamic VLAN Creation**—Displays whether interface dynamic VLAN creation is **Enabled** or **Disabled**.

show mvrp

```
user@switch> show mvrp

Global MVRP configuration
```

```
MVRP status               : Enabled
MVRP dynamic vlan creation: Enabled
MVRP Timers (ms):
 Interface        Join   Leave   LeaveAll
 --------------   ----   -----   --------
 all              200    600     10000
 xe-0/1/1.0       200    600     10000

Interface based configuration:
 Interface        Status     Registration   Dynamic VLAN Creation
 --------------   --------   ------------   ---------------------
 all       Disabled   Normal    Enabled
 xe-0/1/1.0     Enabled    Normal         Enabled
```

show mvrp dynamic-vlan-memberships

Syntax show mvrp dynamic-vlan-memberships

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display all VLANs that have been created dynamically using Multiple VLAN Registration Protocol (MVRP) on the switch.

Required Privilege Level clear

Related Documentation

- **show mvrp on page 1257**
- **show mvrp statistics on page 1260**
- Example: Configuring Automatic VLAN Administration Using MVRP on J-EX Series Switches on page 1115
- Verifying That MVRP Is Working Correctly on page 1162

List of Sample Output **show mvrp dynamic-vlan-memberships on page 1259**

Output Fields Table 167 on page 1259 lists the output fields for the **show mvrp dynamic-vlan-memberships** command. Output fields are listed in the approximate order in which they appear.

Table 167: show mvrp dynamic-vlan-memberships Output Fields

Field Name	Field Description
VLAN Name	The name of the dynamically created VLAN.
Interfaces	The interface or interfaces that are bound to the dynamically created VLAN.

show mvrp dynamic-vlan-memberships

```
user@switch> show mvrp dynamic-vlan-memberships
VLAN Name              Interfaces
-------------------        ----------------
__mvrp_100__               xe-0/1/1.0
                           xe-0/1/0.0
__mvrp_200__               xe-0/1/1.0
                           xe-0/1/0.0
__mvrp_300__               xe-0/1/1.0
```

show mvrp statistics

Syntax show mvrp statistics <interface *interface-name*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display Multiple VLAN Registration Protocol (MVRP) statistics in the form of Multiple Registration Protocol data unit (MRPDU) messages.

Options none—Show MVRP statistics for all interfaces on the switch.

interface *interface-name*—Show MVRP statistics for the specified interface.

Required Privilege Level clear

Related Documentation

- **show mvrp on page 1257**
- **clear mvrp statistics on page 1230**
- Example: Configuring Automatic VLAN Administration Using MVRP on J-EX Series Switches on page 1115
- Verifying That MVRP Is Working Correctly on page 1162

List of Sample Output **show mvrp statistics interface xe-0/1/1.0 on page 1261**

Output Fields Table 168 on page 1260 lists the output fields for the **show mvrp statistics** command. Output fields are listed in the approximate order in which they appear.

Table 168: show mvrp statistics Output Fields

Field Name	Field Description
MRPDU received	Number of MRPDU messages received on the switch.
Invalid PDU received	Number of invalid MRPDU messages received on the switch.
New received	Number of new messages received on the switch.
Join Empty received	Number of MRP Join Empty messages received on the switch.
Join In received	Number of MRP Join In messages received on the switch.
Empty received	Number of MRP Empty messages received on the switch.
In received	Number of MRP In messages received on the switch.
Leave received	Number of MRP Leave messages received on the switch.
LeaveAll received	Number of LeaveAll messages received on the switch.

Table 168: show mvrp statistics Output Fields *(continued)*

Field Name	Field Description
MRPDU transmitted	Number of MRPDU messages transmitted from the switch.
MRPDU transmit failures	Number of MRPDU transmit failures from the switch.
New transmitted	Number of new messages transmitted from the switch.
Join Empty transmitted	Number of Join Empty messages sent from the switch.
Join In transmitted	Number of MRP Join In messages sent from the switch.
Empty transmitted	Number of MRP Empty messages sent from the switch.
In transmitted	Number of MRP In messages sent from the switch.
Leave transmitted	Number of MRP Leave Empty messages sent from the switch.
LeaveAll transmitted	Number of MRP LeaveAll messages sent from the switch.

show mvrp statistics interface xe-0/1/1.0

```
user@switch> show mvrp statistics interface xe-0/1/1.0
MVRP statistics
  MRPDU received                : 3342
  Invalid PDU received          : 0
  New received                  : 2
  Join Empty received           : 1116
  Join In received              : 2219
  Empty received                : 2
  In received                   : 2
  Leave received                : 1
  LeaveAll received             : 1117
  MRPDU transmitted             : 3280
  MRPDU transmit failures       : 0
  New transmitted               : 0
  Join Empty transmitted        : 1114
  Join In transmitted           : 2163
  Empty transmitted             : 1
  In transmitted                : 1
  Leave transmitted             : 1
  LeaveAll transmitted          : 1111
```

show redundant-trunk-group

Syntax **show redundant-trunk-group <group-name *group-name*>**

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display information about redundant trunk groups.

Options group-name *group-name*—Display information about the specified redundant trunk group.

Required Privilege Level view

Related Documentation
- Example: Configuring Redundant Trunk Links for Faster Recovery on page 1101
- Understanding Redundant Trunk Links on J-EX Series Switches on page 1049

List of Sample Output **show redundant-trunk-group group-name Group1 on page 1262**

Output Fields Table 169 on page 1262 lists the output fields for the **show redundant-trunk-group** command. Output fields are listed in the approximate order in which they appear.

Table 169: show redundant-trunk-group Output Fields

Field Name	Field Description
Group Name	Name of the redundant trunk port group.
Interface	Name of an interface belonging to the trunk port group. • **(P)** denotes a primary interface. • **(A)** denotes an active interface. • Lack of **(A)** denotes a blocking interface.
State	Operating state of the interface: **UP** or **DOWN**.
Last Time of Flap	Date and time at which the advertised link became unavailable, and then, available again.
# Flaps	Total number of flaps since the last switch reboot.

show redundant-trunk-group group-name Group1

```
user@switch> show redundant—trunk-group group-name Group1
show redundant-trunk-group group-name Group1

Group Name Interface         State     Last Time of Flap       # Flaps
  Group1   ge-0/0/45.0 (P)     UP      Fri Jan  2 04:10:58      0
           ge-0/0/47.0         UP      Fri Jan  2 04:10:58      0
```

show vlans

Syntax

show vlans
<brief | detail | extensive>
<dot1q-tunneling>
<management-vlan>
<sort-by (name | tag)>
<summary>
<*vlan-name*>
<*vlan-range-name*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display information about VLANs configured on bridged Ethernet interfaces. For interfaces configured to support a VoIP VLAN and a data VLAN, the **show vlans** command displays both tagged and untagged membership for those VLANs.

NOTE: When a series of VLANs is created with the vlan-range statement, such VLAN names are prefixed and suffixed with a double underscore. For example, a series of VLANs using the VLAN range 1–3 and the base VLAN name marketing are displayed as __marketing_1__, __marketing_2__, and __marketing_3__.

NOTE: To display an 802.1X supplicant successfully authenticated in multiple-supplicant mode with dynamic VLAN movement, use the show vlans *vlan-name* extensive operational mode command, where *vlan-name* is the dynamic VLAN.

Options none—Display information for all VLANs. VLAN information is displayed by VLAN name in ascending order.

brief | detail | extensive—(Optional) Display the specified level of output.

dot1q-tunneling—(Optional) Display VLANs with the Q-in-Q tunneling feature enabled.

management-vlan—(Optional) Display management VLANs.

sort-by (name | tag)—(Optional) Display VLANs in ascending order of VLAN IDs or VLAN names.

summary—(Optional) Display the total number of VLANs and counts of VLANs by type—for example, the number of dynamic, 802.1Q, and Q-in-Q tunneled VLANs.

vlan-name—(Optional) Display information for the specified VLAN.

vlan-range-name—(Optional) Display information for the specified VLAN range. To see information for all members of the VLAN range, specify the base VLAN name—for

example, **employee** for a VLAN range that includes **__employee_1__** through **__employee_10__**.

Required Privilege Level view

Related Documentation

- **show ethernet-switching interfaces on page 997**
- Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch on page 1063
- Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches on page 1070
- Example: Configure Automatic VLAN Administration Using GVRP on page 1087
- Example: Configuring a Private VLAN on a J-EX Series Switch on page 1107
- Example: Setting Up Q-in-Q Tunneling on J-EX Series Switches on page 1105
- Understanding Bridging and VLANs on J-EX Series Switches on page 1041

List of Sample Output

show vlans on page 1266
show vlans brief on page 1267
show vlans detail on page 1267
show vlans extensive (MAC-based) on page 1268
show vlans extensive (Port-based) on page 1268
show vlans sort-by tag on page 1269
show vlans sort-by name on page 1270
show vlans employee (vlan-range-name) on page 1270
show vlans summary on page 1271

Output Fields Table 170 on page 1264 lists the output fields for the **show vlans** command. Output fields are listed in the approximate order in which they appear.

Table 170: show vlans Output Fields

Field Name	Field Description	Level of Output
Name	Name of a VLAN.	none, **brief**
Tag	The 802.1Q tag applied to this VLAN. If **none** is displayed, no tag is applied.	All levels
Interfaces	Interface associated with learned MAC addresses or all-members (flood entry). An asterisk (*) beside the interface indicates that the interface is **UP**.	All levels
Address	The IP address.	none, **brief**
Ports Active / Total	The number of interfaces associated with a VLAN. The **Active** column indicates interfaces that are **UP**, and the **Total** column indicates interfaces that are active and inactive.	**brief**
VLAN	Name of a VLAN.	**detail**, **extensive**
Admin state	Indicates whether the physical link is operational and can pass packets.	**detail**, **extensive**

Table 170: show vlans Output Fields *(continued)*

Field Name	Field Description	Level of Output
Dot1q Tunneling Status	Indicates whether Q-in-Q tunneling is enabled.	**detail**, **extensive**
MAC learning Status	Indicates whether MAC learning is disabled.	**detail**, **extensive**
Description	A description for the VLAN.	**detail**,**extensive**
Primary IP	Primary IP address associated with a VLAN.	**detail**
Number of interfaces	The number of interfaces associated with a VLAN. Both the total number of interfaces and the number of active interfaces associated with a VLAN are displayed.	**detail**, **extensive**
STP	The spanning tree associated with a VLAN.	**detail**, **extensive**
RTG	The redundant trunk group associated with a VLAN.	**detail**, **extensive**
Tagged interfaces	The tagged interfaces to which a VLAN is associated.	**detail**, **extensive**
Untagged interfaces	The untagged interfaces to which a VLAN is associated.	**detail**. **extensive**
Customer VLAN Ranges	Lists the customer VLAN (C-VLAN) ranges associated with this service VLAN (S-VLAN).	**extensive**
Private VLAN Mode	The private VLAN mode for this VLAN. Values are **Primary**, **Isolated**, and **Community**.	**extensive**
Primary VLAN	The primary VLAN tag for this secondary VLAN.	**extensive**
Interrnal Index	VLAN index internal to the Junos OS.	**extensive**
Origin	The manner in which the VLAN was created. Values are **static** and **learn**.	**extensive**
Protocol	Port-based VLAN or MAC-based VLAN. MAC-based protocol is displayed when VLAN assignment is done either statically or dynamically through 802.1X.	**extensive**
Mac aging time	The MAC aging timer.	**extensive**
IP addresses	IP address associated with a VLAN.	**extensive**
Number of MAC entries	For MAC-based VLANs created either statically or dynamically, the MAC addresses associated with an interface.	**extensive**
Secondary VLANs	The secondary VLANs associated with a primary VLAN.	**extensive**
Isolated VLANs	The isolated VLANs associated with a primary VLAN.	**extensive**
Community VLANs	The community VLANs associated with a primary VLAN.	**extensive**

Table 170: show vlans Output Fields *(continued)*

Field Name	Field Description	Level of Output
VLANs summary	VLAN counts: • **Total**—Total number of VLANs on the switch. • **Configured VLANs**—Number of VLANs that are based on user-configured settings. • **Internal VLANs**—Number of VLANs created by the system with no explicit configuration or protocol—for example, the **default** VLAN and the VLAN created when a trunk interface is not configured with native VLAN membership. • **Temporary VLANs**—Number of VLANs from the previous configuration that the system retains for a limited time after restart. Temporary VLANs are converted into one of the other types of VLAN, or are removed from the system if the current configuration does not require them.	All levels

Dot1q VLANs summary	802.1Q VLAN counts:
	• **Total**—Total number of 802.1Q VLANs on the switch.
	• **Tagged VLANs**—Number of tagged 802.1Q VLANs.
	• **Untagged VLANs**—Number of untagged 802.1Q VLANs.
	• **Private VLAN**—Counts of the following kinds of 802.1Q private VLANs (PVLANs):
	  ◦ **Primary VLANs**—Number of primary forwarding private VLANs.
	  ◦ **Community VLANs**—Number of secondary transporting and forwarding private VLANs.
	  ◦ **Isolated VLANs**—Number of secondary receiving and forwarding private VLANs.
		All levels

Dot1q Tunneled VLANs summary	Q-in-Q VLAN counts: • **Total**—Total number of Q-in-Q VLANs on the switch. • **Private VLAN**—Counts of primary, community, and isolated Q-in-Q private VLANs (PVLANs).	All levels
Dynamic VLANs	Counts of VLANs assigned or created dynamically by a protocol: • **Total**—Total number of dynamic VLANs on the switch. • **Dot1x**—Number of 802.1X VLANs authenticated and assigned when the switch learns the MAC address of a supplicant host from a packet's source MAC address. • **MVRP**—Number of VLANs created by the Multiple VLAN Registration Protocol (MVRP).	All levels

show vlans

```
user@switch> show vlans

Name            Tag     Interfaces
default         None
                        ge-0/0/34.0, ge-0/0/33.0, ge-0/0/32.0, ge-0/0/31.0,
                        ge-0/0/30.0, ge-0/0/29.0, ge-0/0/28.0, ge-0/0/27.0,
                        ge-0/0/26.0, ge-0/0/25.0, ge-0/0/19.0, ge-0/0/18.0,
                        ge-0/0/17.0, ge-0/0/16.0, ge-0/0/15.0, ge-0/0/14.0,
                        ge-0/0/13.0, ge-0/0/11.0, ge-0/0/9.0, ge-0/0/8.0,
```

```
                                ge-0/0/3.0, ge-0/0/2.0, ge-0/0/1.0
v0001           1
                                ge-0/0/24.0, ge-0/0/23.0, ge-0/0/22.0, ge-0/0/21.0
v0002           2
                                None
v0003           3
                                None
v0004           4
                                None
v0005           5
                                None
```

show vlans brief

```
user@switch> show vlans brief
                                                Ports
Name            Tag     Address                 Active/Total
default         None                            0/23
v0001           1                               0/4
v0002           2                               0/0
v0003           3                               0/0
v0004           4                               0/0
v0005           5                               0/0
v0006           6                               0/0
v0007           7                               0/0
v0008           8                               0/0
v0009           9                               0/0
v0010           10                              0/2
v0011           11                              0/0
v0012           12                              0/0
v0013           13                              0/0
v0014           14                              0/0
v0015           15                              0/0
v0016           16                              0/0
```

show vlans detail

```
user@switch> show vlans detail
VLAN: default, Tag: Untagged, Admin state: Enabled
  Description: None
  Primary IP: None, Number of interfaces: 23 (Active = 0)
  STP: None, RTG: None
  Untagged interfaces: ge-0/0/34.0, ge-0/0/33.0, ge-0/0/32.0, ge-0/0/31.0,
  ge-0/0/30.0, ge-0/0/29.0, ge-0/0/28.0, ge-0/0/27.0, ge-0/0/26.0,
  ge-0/0/25.0, ge-0/0/19.0, ge-0/0/18.0, ge-0/0/17.0, ge-0/0/16.0,
  ge-0/0/15.0, ge-0/0/14.0, ge-0/0/13.0, ge-0/0/11.0, ge-0/0/9.0, ge-0/0/8.0,
  ge-0/0/3.0, ge-0/0/2.0, ge-0/0/1.0,
  Tagged interfaces: None

VLAN: v0001, Tag: 802.1Q Tag 1, Admin state: Enabled
  Description: None
  Primary IP: None, Number of interfaces: 4 (Active = 0)
  Dot1q Tunneling Status: Enabled
  STP: None, RTG: None
  Untagged interfaces: None
  Tagged interfaces: ge-0/0/24.0, ge-0/0/23.0, ge-0/0/22.0, ge-0/0/21.0,

VLAN: v0002, Tag: 802.1Q Tag 2, Admin state: Enabled
  Description: None
  Primary IP: None, Number of interfaces: 0 (Active = 0)
  STP: None, RTG: None
  Untagged interfaces: None
  Tagged interfaces: None

VLAN: v0003, Tag: 802.1Q Tag 3, Admin state: Enabled
```

```
  Description: None
  Primary IP: None, Number of interfaces: 0 (Active = 0)
  STP: None, RTG: None
  Untagged interfaces: None
  Tagged interfaces: None

VLAN: vlan4000, 802.1Q Tag: Untagged, Admin State: Enabled
MAC learning Status: Disabled
Number of interfaces: 0 (Active = 0)
```

show vlans extensive (MAC-based)

```
user@switch> show vlans extensive
VLAN: default, Created at: Thu May 15 13:43:09 2008
Internal index: 3, Admin State: Enabled, Origin: Static
Protocol: Port Mode, Mac aging time: 300 seconds
Number of interfaces: Tagged 0 (Active = 0), Untagged  2 (Active = 2)
      ge-0/0/0.0*, untagged, access
      ge-0/0/14.0*, untagged, access

VLAN: vlan_dyn, Created at: Thu May 15 13:43:09 2008
Internal index: 4, Admin State: Enabled, Origin: Static
Protocol: Port Mode
Number of interfaces: Tagged 0 (Active = 0), Untagged  0 (Active = 0)
Protocol: MAC Based
Number of MAC entries: 6
      ge-0/0/0.0*
                  00:00:00:00:00:02 (untagged)
                  00:00:00:00:00:03 (untagged)
                  00:00:00:00:00:04 (untagged)
                  00:00:00:00:00:05 (untagged)
                  00:00:00:00:00:06 (untagged)
                  00:00:00:00:00:07 (untagged)
```

show vlans extensive (Port-based)

```
user@switch> show vlans extensive
VLAN: default, created at Mon Feb  4 12:13:47 2008
  Tag: None, Internal index: 0, Admin state: Enabled, Origin: static
  Description: None
  Dot1q Tunneling Status: Enabled
  Customer VLAN ranges:
                    1-4100
  Private VLAN Mode: Primary
  Protocol: Port based, Layer 3 interface: None
  IP addresses: None
  STP: None, RTG: None.
  Number of interfaces: Tagged 0 (Active = 0), Untagged 23 (Active = 0)
    ge-0/0/34.0 (untagged, access)
    ge-0/0/33.0 (untagged, access)
    ge-0/0/32.0 (untagged, access)
    ge-0/0/31.0 (untagged, access)
    ge-0/0/30.0 (untagged, access)
    ge-0/0/29.0 (untagged, access)
    ge-0/0/28.0 (untagged, access)
    ge-0/0/27.0 (untagged, access)
    ge-0/0/26.0 (untagged, access)
    ge-0/0/25.0 (untagged, access)
    ge-0/0/19.0 (untagged, access)
    ge-0/0/18.0 (untagged, access)
    ge-0/0/17.0 (untagged, access)
    ge-0/0/16.0 (untagged, access)
    ge-0/0/15.0 (untagged, access)
    ge-0/0/14.0 (untagged, access)
    ge-0/0/13.0 (untagged, access)
```

```
    ge-0/0/11.0 (untagged, access)
    ge-0/0/9.0 (untagged, access)
    ge-0/0/8.0 (untagged, access)
    ge-0/0/3.0 (untagged, access)
    ge-0/0/2.0 (untagged, access)
    ge-0/0/1.0 (untagged, access)

Secondary VLANs: Isolated 1, Community  1
  Isolated VLANs :
      __pvlan_pvlan_ge-0/0/3.0__
  Community VLANs :
      comm1

VLAN: v0001, created at Mon Feb  4 12:13:47 2008
  Tag: 1, Internal index: 1, Admin state: Enabled, Origin: static
  Description: None
  Protocol: Port based, Layer 3 interface: None
  IP addresses: None
  STP: None, RTG: None.
  Number of interfaces: Tagged 4 (Active = 0), Untagged 0 (Active = 0)
    ge-0/0/24.0 (tagged, trunk)
    ge-0/0/23.0 (tagged, trunk)
    ge-0/0/22.0 (tagged, trunk)
    ge-0/0/21.0 (tagged, trunk)

VLAN: v0002, created at Mon Feb  4 12:13:47 2008
  Tag: 2, Internal index: 2, Admin state: Enabled, Origin: static
  Description: None
  Protocol: Port based, Layer 3 interface: None
  IP addresses: None
  STP: None, RTG: None.
  Number of interfaces: Tagged 0 (Active = 0), Untagged 0 (Active = 0)
    None

VLAN: v0003, created at Mon Feb  4 12:13:47 2008
  Tag: 3, Internal index: 3, Admin state: Enabled, Origin: static
  Description: None
  Protocol: Port based, Layer 3 interface: None
  IP addresses: None
  STP: None, RTG: None.
  Number of interfaces: Tagged 0 (Active = 0), Untagged 0 (Active = 0)
    None
```

show vlans sort-by tag

```
user@switch> show vlans sort-by tag
Name           Tag     Interfaces
default                None
__vlan-x_1__   1
                       None
__vlan-x_2__   2
                       None
__vlan-x_3__   3
                       None
__vlan-x_4__   4
                       None
__vlan-x_5__   5
                       None
__vlan-x_6__   6
                       None
__vlan-x_7__   7
                       None
__vlan-x_8__   8
```

```
                                    None
__vlan-x_9__    9
                                    None
__vlan-x_10__   10
                                    None
__vlan-x_11__   11
                                    None
__vlan-x_12__   12
                                    None
__vlan-x_13__   13
                                    None
__vlan-x_14__   14
                                    None
__vlan-x_15__   15
                                    None
__vlan-x_16__   16
                                    None
__vlan-x_17__   17
                                    None
__vlan-x_18__   18
                                    None
__vlan-x_19__   19
                                    None
__vlan-x_20__   20
                                    None
```

show vlans sort-by name

```
user@switch> show vlans sort-by name
Name              Tag   Interfaces

__employee_120__ 120
                        ge-0/0/22.0*
__employee_121__ 121
                        ge-0/0/22.0*
__employee_122__ 122
                        ge-0/0/22.0*
__employee_123__ 123
                        ge-0/0/22.0*
__employee_124__ 124
                        ge-0/0/22.0*
__employee_125__ 125
                        ge-0/0/22.0*
__employee_126__ 126
                        ge-0/0/22.0*
__employee_127__ 127
                        ge-0/0/22.0*
__employee_128__ 128
                        ge-0/0/22.0*
__employee_129__ 129
                        ge-0/0/22.0*
__employee_130__ 130
                        ge-0/0/22.0*
```

show vlans employee (vlan-range-name)

```
user@switch> show vlans employee
Name              Tag   Interfaces

__employee_120__ 120
                        ge-0/0/22.0*
__employee_121__ 121
                        ge-0/0/22.0*
```

```
__employee_122__ 122
                        ge-0/0/22.0*
__employee_123__ 123
                        ge-0/0/22.0*
__employee_124__ 124
                        ge-0/0/22.0*
__employee_125__ 125
                        ge-0/0/22.0*
__employee_126__ 126
                        ge-0/0/22.0*
__employee_127__ 127
                        ge-0/0/22.0*
__employee_128__ 128
                        ge-0/0/22.0*
__employee_129__ 129
                        ge-0/0/22.0*
__employee_130__ 130
                        ge-0/0/22.0*
```

show vlans summary

```
user@switch> show vlans summary
VLANs summary:
         Total: 8,   Configured VLANs: 5
         Internal VLANs: 1,  Temporary VLANs: 0

     Dot1q VLANs summary:
         Total: 8,  Tagged VLANs: 2, Untagged VLANs: 6
         Private VLAN:
             Primary VLANs: 2,  Community VLANs: 2, Isolated VLANs: 3

     Dot1q Tunneled VLANs summary:
         Total: 0
         Private VLAN:
             Primary VLANs: 0,  Community VLANs: 0, Isolated VLANs: 0

     Dynamic VLANs:
        Total: 2,  Dot1x: 2, MVRP: 0
```

PART 14

Spanning-Tree Protocols

CHAPTER 64

Spanning-Tree Protocols—Overview

Understanding STP for J-EX Series Switches

J-EX Series Switches provide Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), Multiple Spanning Tree Protocol (MSTP), and VLAN Spanning Tree Protocol (VSTP). The default spanning-tree protocol for J-EX Series switches is RSTP. RSTP provides faster convergence times than STP. However, some legacy networks require the slower convergence times of basic STP.

If your network includes 802.1D 1998 bridges, you can remove RSTP and explicitly configure STP. See “Configuring STP (CLI Procedure)” on page 1336. When you explicitly configure STP, the J-EX Series switches use the IEEE 802.1D 2004 specification, force version 0. This configuration runs a version of RSTP that is compatible with the classic, basic STP. If you use VLANs, you should enable VSTP and use it on your network. See “Understanding VSTP for J-EX Series Switches” on page 1281.

You can use the same operational commands (**show spanning-tree bridge** and **show spanning-tree interface**) to check the status of your spanning-tree configuration, regardless of which spanning-tree protocol has been configured.

STP uses bridge protocol data unit (BPDU) packets to exchange information with other switches. BPDUs send hello packets out at regular intervals to exchange information across bridges and detect loops in a network topology. There are two types of BPDUs:

- Configuration BPDUs: Contain configuration information about the transmitting switch and its ports, including switch and port MAC addresses, switch priority, port priority, and port cost.
- Topology Change Notification (TCN) BPDUs: When a bridge needs to signal a topology change, it starts to send TCNs on its root port. The designated bridge receives the TCN, acknowledges it, and generates another one for its own root port. The process continues until the TCN reaches the root bridge.

STP uses the information provided by the BPDUs to elect a root bridge, identify root ports for each switch, identify designated ports for each physical LAN segment, and prune specific redundant links to create a loop-free tree topology. All leaf devices calculate the best path to the root device and place their ports in blocking or forwarding states based on the best path to the root. The resulting tree topology provides a single active Layer 2 data path between any two end stations.

Related Documentation

- Understanding MSTP for J-EX Series Switches on page 1277
- Understanding RSTP for J-EX Series Switches on page 1276
- Understanding VSTP for J-EX Series Switches on page 1281
- Understanding Layer 2 Protocol Tunneling on J-EX Series Switches on page 1056

Understanding RSTP for J-EX Series Switches

J-EX Series Switches use Rapid Spanning Tree Protocol (RSTP) to provide better reconvergence time than that provided by the base Spanning Tree Protocol (STP). RSTP identifies certain links as point to point. When a point-to-point link fails, the alternate link can transition to the forwarding state.

Although STP provides basic loop prevention functionality, it does not provide fast network convergence when there are topology changes. STP's process to determine network state transitions is slower than RSTP's because it is timer-based. A device must reinitialize every time a topology change occurs. The device must start in the listening state and transition to the learning state and eventually to a forwarding or blocking state. When default values are used for the maximum age (20 seconds) and forward delay (15 seconds), it takes 50 seconds for the device to converge. RSTP converges faster because it uses a handshake mechanism based on point-to-point links instead of the timer-based process used by STP.

For networks with VLANs, you can use VLAN Spanning Tree Protocol (VSTP) to run one or more STP or RSTP instances for each VLAN on which VSTP is enabled. VSTP takes the paths of each VLAN into account when calculating routes.. VSTP uses RSTP instances by default.

An RSTP domain running on a switch has the following components:

- A *root port*, which is the “best path” to the root device.
- A *designated port*, which indicates that the switch is the designated bridge for the other switch connecting to this port.

- An *alternate port*, which provides an alternate root port.
- A *backup port*, which provides an alternate designated port.

Port assignments change through messages exchanged throughout the domain. An RSTP device generates configuration messages once per every hello time interval. If an RSTP device does not receive a configuration message from its neighbor after an interval of three hello times, it determines that the connection with the neighbor is lost. When a *root port* or a *designated port* fails on a device, the device generates a configuration message with the proposal bit set. Once its neighbor device receives this message, it verifies that this configuration message is better than the one saved for that port and then it starts a *synchronizing* operation to ensure that all of its ports are in sync with the new information.

Similar waves of proposal agreement handshake messages propagate toward the leaves of the network, quickly restoring the connectivity after a topology change (in a well-designed network that uses RSTP, network convergence can take as little as 0.5 seconds). If a device does not receive an agreement to a proposal message it has sent, it returns to the original IEEE 802.D convention.

RSTP was originally defined in the IEEE 802.1w draft specification and later incorporated into the IEEE 802.1D-2004 specification.

VSTP and RSTP can be configured concurrently. You can selectively configure up to 253 VLANs using VSTP; the remaining VLANs will be configured using RSTP. VSTP and RSTP are the only spanning-tree protocols that can be configured concurrently on the switch.

Related Documentation

- Understanding STP for J-EX Series Switches on page 1275
- Understanding MSTP for J-EX Series Switches on page 1277
- Understanding VSTP for J-EX Series Switches on page 1281
- Understanding Layer 2 Protocol Tunneling on J-EX Series Switches on page 1056
- Example: Configuring Faster Convergence and Improving Network Stability with RSTP on J-EX Series Switches on page 1283

Understanding MSTP for J-EX Series Switches

Although RSTP provides faster convergence time than STP, it still does not solve a problem inherent in STP: All VLANs within a LAN must share the same spanning tree. To solve this problem, J-EX Series Switches use Multiple Spanning Tree Protocol (MSTP) to create a loop-free topology in networks with multiple spanning-tree regions.

An MSTP region allows a group of bridges to be modeled as a single bridge. An MSTP region contains multiple spanning tree instances (MSTIs). MSTIs provide different paths for different VLANs. This functionality facilitates better load sharing across redundant links.

MSTP region can support up to 64 MSTIs and each instance can support anywhere from 1 through 4094 vlans.

MSTP was originally defined in the IEEE 802.1s draft specification and later incorporated into the IEEE 802.1Q-2003 specification.

Related Documentation

- Understanding STP for J-EX Series Switches on page 1275
- Understanding RSTP for J-EX Series Switches on page 1276
- Understanding Layer 2 Protocol Tunneling on J-EX Series Switches on page 1056
- Example: Configuring Network Regions for VLANs with MSTP on J-EX Series Switches on page 1297

Understanding BPDU Protection for STP, RSTP, and MSTP on J-EX Series Switches

J-EX Series Switches provide Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), and Multiple Spanning Tree Protocol (MSTP). BPDU protection can help prevent STP misconfigurations that can lead to network outages.

A loop-free network is supported through the exchange of a special type of frame called bridge protocol data unit (BPDU). Receipt of BPDUs on certain interfaces in an STP, RSTP, VSTP, or MSTP topology, however, can lead to network outages. Enable BPDU protection on those interfaces to prevent these outages.

Peer STP applications running on the switch interfaces use BPDUs to communicate. Ultimately, the exchange of BPDUs determines which interfaces block traffic and which interfaces become root ports and forward traffic.

However, a user bridge application running on a PC can also generate BPDUs. If these BPDUs are picked up by STP applications running on the switch, they can trigger STP miscalculations, and those miscalculations can lead to network outages.

Enable BPDU protection on switch interfaces connected to user devices or on interfaces on which no BPDUs are expected, such as edge ports. If BPDUs are received on a protected interface, the interface is disabled and stops forwarding frames.

Not only can you configure BPDU protection on a switch with a spanning tree, but also on a switch without a spanning tree. This type of topology typically consists of a non-STP switch connected to an STP switch through a trunk interface.

To configure BPDU protection on a switch with a spanning tree, include the **bpdu-block-on-edge** statement at the [**edit protocols** (**stp** | **mstp** | **rstp**)] hierarchy level. To configure BPDU protection on a switch without a spanning tree, include the **bpdu-block** statement at the [**edit ethernet-switching-options interface** *interface-name*] hierarchy level.

After the misconfiguration that triggered the BPDUs being sent to an interface is fixed in the topology, the interface can be unblocked in one of two ways:

- If the **disable-timeout** statement has been included in the BPDU configuration, the interface automatically returns to service after the timer expires.
- Use the operational mode command **clear ethernet-switching bpdu-error**.

Disabling the BPDU protection configuration does not unblock the interface.

Related Documentation

- Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations on J-EX Series Switches on page 1317
- Example: Configuring BPDU Protection on non-STP Interfaces to Prevent STP Miscalculations on J-EX Series Switches on page 1321
- Understanding Loop Protection for STP, RSTP, VSTP, and MSTP on J-EX Series Switches on page 1279
- Understanding Root Protection for STP, RSTP, VSTP, and MSTP on J-EX Series Switches on page 1280
- Understanding MSTP for J-EX Series Switches on page 1277
- Understanding RSTP for J-EX Series Switches on page 1276
- Understanding STP for J-EX Series Switches on page 1275
- Understanding VSTP for J-EX Series Switches on page 1281

Understanding Loop Protection for STP, RSTP, VSTP, and MSTP on J-EX Series Switches

J-EX Series Switches provide Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), and Multiple Spanning Tree Protocol (MSTP). Loop protection increases the efficiency of STP, RSTP, and MSTP by preventing ports from moving into a forwarding state that would result in a loop opening up in the network.

A loop-free network in spanning-tree topologies is supported through the exchange of a special type of frame called bridge protocol data unit (BPDU). Peer STP applications running on the switch interfaces use BPDUs to communicate. Ultimately, the exchange of BPDUs determines which interfaces block traffic (preventing loops) and which interfaces become root ports and forward traffic.

However, a blocking interface can transition to the forwarding state in error if the interface stops receiving BPDUs from its designated port on the segment. Such a transition error can occur when there is a hardware error on the switch or software configuration error between the switch and its neighbor.

When loop protection is enabled, the spanning-tree topology detects root ports and blocked ports and makes sure both keep receiving BPDUs. If a loop-protection-enabled interface stops receiving BPDUs from its designated port, it reacts as it would react to a problem with the physical connection on this interface. It doesn't transition the interface to a forwarding state, but instead transitions it to a loop-inconsistent state. The interface recovers and then it transitions back to the spanning-tree blocking state as soon as it receives a BPDU.

We recommend that you enable loop protection on all switch interfaces that have a chance of becoming root or designated ports. Loop protection is most effective when

enabled in the entire switched network. When you enable loop protection, you must configure at least one action (**alarm**, **block**, or both).

An interface can be configured for either loop protection or root protection, but not for both.

Related Documentation

- Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on J-EX Series Switches on page 1325
- Understanding Root Protection for STP, RSTP, VSTP, and MSTP on J-EX Series Switches on page 1280
- Understanding BPDU Protection for STP, RSTP, and MSTP on J-EX Series Switches on page 1278
- Understanding MSTP for J-EX Series Switches on page 1277
- Understanding RSTP for J-EX Series Switches on page 1276
- Understanding STP for J-EX Series Switches on page 1275
- Understanding VSTP for J-EX Series Switches on page 1281

Understanding Root Protection for STP, RSTP, VSTP, and MSTP on J-EX Series Switches

J-EX Series Switches provide Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), and Multiple Spanning Tree Protocol (MSTP). A loop-free network is supported through the exchange of a special type of frame called bridge protocol data unit (BPDU). Peer STP applications running on the switch interfaces use BPDUs to communicate. Ultimately, the exchange of BPDUs determines which interfaces block traffic and which interfaces become root ports and forward traffic.

However, a root port elected through this process has the possibility of being wrongly elected. A user bridge application running on a PC can generate BPDUs, too, and interfere with root port election. Root protection allows network administrators to manually enforce the root bridge placement in the network.

Enable root protection on interfaces that should not receive superior BPDUs from the root bridge and should not be elected as the root port. These interfaces become designated ports and are typically located on an administrative boundary. If the bridge receives superior STP BPDUs on a port that has root protection enabled, that port transitions to a root-prevented STP state (inconsistency state) and the interface is blocked. This blocking prevents a bridge that should not be the root bridge from being elected the root bridge. After the bridge stops receiving superior STP BPDUs on the interface with root protection, the interface returns to a listening state, followed by a learning state, and ultimately back to a forwarding state. Recovery back to the forwarding state is automatic.

When root protection is enabled on an interface, it is enabled for all the STP instances on that interface. The interface is blocked only for instances for which it receives superior BPDUs. Otherwise, it participates in the spanning-tree topology.

An interface can be configured for either root protection or loop protection, but not for both.

Related Documentation

- Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees on J-EX Series Switches on page 1329
- Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on J-EX Series Switches on page 1325
- Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations on J-EX Series Switches on page 1317
- Example: Configuring BPDU Protection on non-STP Interfaces to Prevent STP Miscalculations on J-EX Series Switches on page 1321
- Understanding MSTP for J-EX Series Switches on page 1277
- Understanding RSTP for J-EX Series Switches on page 1276
- Understanding STP for J-EX Series Switches on page 1275
- Understanding VSTP for J-EX Series Switches on page 1281

Understanding VSTP for J-EX Series Switches

VLAN Spanning Tree Protocol (VSTP) allows J-EX Series Switches to run one or more Spanning Tree Protocol (STP) or Rapid Spanning Tree Protocol (RSTP) instances for each VLAN on which VSTP is enabled. For networks with multiple VLANs, VSTP improves intelligent tree spanning by defining best paths within the VLANs instead of within the entire network.

You can configure VSTP for a maximum of 253 VLANs. If you need to run a spanning-tree protocol on more than 253 VLANs, you must configure VSTP and RSTP concurrently. When VSTP and RSTP are configured, up to 253 VLANs can use VSTP and the remaining VLANs use RSTP. You can selectively configure which VLANs use VSTP when VSTP and RSTP are configured.

VSTP and RSTP are the only spanning-tree protocols that can be configured concurrently on the switch.

NOTE: We recommend that you enable VSTP on all VLANs that could receive VSTP bridge protocol data units (BPDUs).

Related Documentation

- Understanding STP for J-EX Series Switches on page 1275
- Understanding RSTP for J-EX Series Switches on page 1276
- Understanding Layer 2 Protocol Tunneling on J-EX Series Switches on page 1056
- Configuring VLAN Spanning Tree Protocol (CLI Procedure) on page 1340

CHAPTER 65

Examples of Spanning-Tree Protocols Configuration

Example: Configuring Faster Convergence and Improving Network Stability with RSTP on J-EX Series Switches

J-EX Series switches use Rapid Spanning Tree Protocol (RSTP) to provide a loop-free topology. RSTP identifies certain links as point to point. When a point-to-point link fails, the alternate link can transition to the forwarding state. RSTP provides better reconvergence time than original STP because it uses protocol handshake messages rather than fixed timeouts. Eliminating the need to wait for timers to expire makes RSTP more efficient than STP.

This example describes how to configure RSTP on four J-EX Series switches:

Requirements

This example uses the following hardware and software components:

- Four J-EX Series switches

Before you configure the switches for RSTP, be sure you have:

- Installed the four switches. See “Connecting and Configuring a J-EX Series Switch (J-Web Procedure)” on page 163.
- Performed the initial software configuration on all switches. See Installing and Connecting a J-EX4200 Switch.

Overview and Topology

In this example, four J-EX Series switches are connected in the topology displayed in Figure 33 on page 1284 to create a loop-free topology.

Figure 33: Network Topology for RSTP

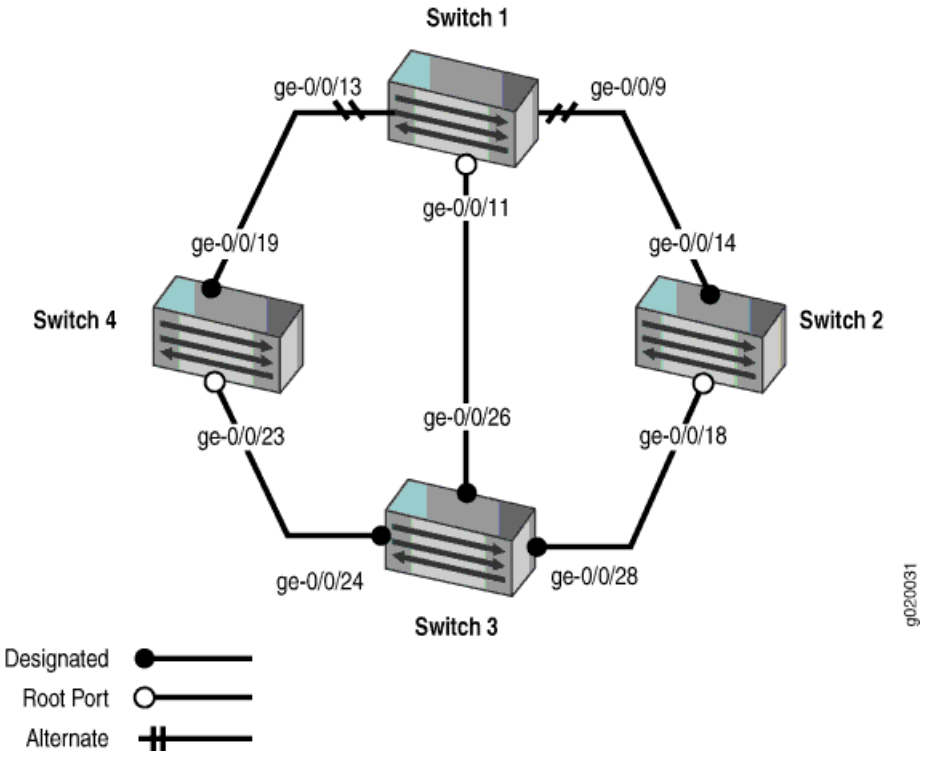

The interfaces shown in Table 171 on page 1285 will be configured for RSTP.

NOTE: You can configure RSTP on logical or physical interfaces. This example shows RSTP configured on logical interfaces.

Table 171: Components of the Topology for Configuring RSTP on J-EX Series Switches

Property	Settings
Switch 1	The following ports on Switch 1 are connected in this way: • **ge-0/0/9** is connected to Switch 2 • **ge-0/0/13** is connected to Switch 4 • **ge-0/0/11** is connected to Switch 3
Switch 2	The following ports on Switch 2 are connected in this way: • **ge-0/0/14** is connected to Switch 1 • **ge-0/0/18** is connected to Switch 3
Switch 3	The following ports on Switch 3 are connected in this way: • **ge-0/0/26** is connected to Switch 1 • **ge-0/0/28** is connected to Switch 2 • **ge-0/0/24** is connected to Switch 4
Switch 4	The following ports on Switch 4 are connected in this way: • **ge-0/0/19** is connected to Switch 1 • **ge-0/0/23** is connected to Switch 3
VLAN names and tag IDs	**voice-vlan**, tag **10** **employee-vlan**, tag **20** **guest-vlan**, tag **30** **camera-vlan**, tag **40**

This configuration example creates a loop-free topology between four J-EX Series switches using RSTP.

An RSTP topology contains ports that have specific roles:

- The root port is responsible for forwarding data to the root bridge.
- The alternate port is a standby port for the root port. When a root port goes down, the alternate port becomes the active root port.
- The designated port forwards data to the downstream network segment or device.
- The backup port is a backup port for the designated port. When a designated port goes down, the backup port becomes the active designated port and starts forwarding data.

NOTE: You also can create a loop-free topology between the aggregation layer and the distribution layer using redundant trunk links. For more information about configuring redundant trunk links, see “Example: Configuring Redundant Trunk Links for Faster Recovery” on page 1101.

Configuring RSTP on Switch 1

To configure RSTP on Switch 1, perform these tasks:

CLI Quick Configuration

To quickly configure interfaces and RSTP on Switch 1, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans voice-vlan description “Voice VLAN”
set vlans voice-vlan vlan-id 10
set vlans employee-vlan description “Employee VLAN”
set vlans employee-vlan vlan-id 20
set vlans guest-vlan description “Guest VLAN”
set vlans guest-vlan vlan-id 30
set vlans camera-vlan description “Camera VLAN”
set vlans camera-vlan vlan-id 40
set interfaces ge-0/0/13 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge–0/0/9 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge–0/0/11 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge–0/0/13 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/9 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/11 unit 0 family ethernet-switching port-mode trunk
set protocols rstp bridge-priority 16k
set protocols rstp interface ge-0/0/13.0 cost 1000
set protocols rstp interface ge-0/0/13.0 mode point-to-point
set protocols rstp interface ge-0/0/9.0 cost 1000
set protocols rstp interface ge-0/0/9.0 mode point-to-point
set protocols rstp interface ge-0/0/11.0 cost 1000
set protocols rstp interface ge-0/0/11.0 mode point-to-point
```

Step-by-Step Procedure

To configure interfaces and RSTP on Switch 1:

1. Configure the VLANs **voice-vlan**, **employee-vlan**, **guest-vlan**, and **camera-vlan**:

```
[edit vlans]
user@switch1# set voice-vlan description “Voice VLAN”
user@switch1# set voice-vlan vlan-id 10
user@switch1# set employee-vlan description “Employee VLAN”
user@switch1# set employee-vlan vlan-id 20
user@switch1# set guest-vlan description “Guest VLAN”
user@switch1# set guest-vlan vlan-id 30
user@switch1# set camera-vlan description “Camera VLAN”
user@switch1# set guest-vlan vlan-id 40
```

2. Configure the VLANs on the interfaces, including support for the Ethernet Switching protocol:

```
[edit interfaces]
user@switch1# set ge-0/0/13 unit 0 family ethernet-switching vlan members [10 20
30 40]
user@switch1# set ge-0/0/9 unit 0 family ethernet-switching vlan members [10 20 30
40]
user@switch1# set ge-0/0/11 unit 0 family ethernet-switching vlan members [10 20
30 40]
```

3. Configure the port mode for the interfaces:

```
[edit interfaces]
user@switch1# set ge-0/0/13 unit 0 family ethernet-switching port-mode trunk
user@switch1# set ge-0/0/9 unit 0 family ethernet-switching port-mode trunk
```

```
user@switch1# set ge-0/0/11 unit 0 family ethernet-switching port-mode trunk
```

4. Configure RSTP on the switch:

```
[edit protocols]
user@switch1# rstp bridge-priority 16k
user@switch1# rstp interface ge-0/0/13.0 cost 1000
user@switch1# rstp interface ge-0/0/13.0 mode point-to-point
user@switch1# rstp interface ge-0/0/9.0 cost 1000
user@switch1# rstp interface ge-0/0/9.0 mode point-to-point
user@switch1# rstp interface ge-0/0/11.0 cost 1000
user@switch1# rstp interface ge-0/0/11.0 mode point-to-point
```

Results Check the results of the configuration:

```
user@switch1> show configuration
interfaces {
    ge-0/0/13 {
        unit 0 {
            family ethernet-switching {
                port-mode trunk;
                vlan {
                    members [10 20 30 40];
                }
            }
        }
    }
    ge-0/0/9 {
        unit 0 {
            family ethernet-switching {
                port-mode trunk;
                vlan {
                    members [10 20 30 40];
                }
            }
        }
    }
    ge-0/0/11 {
        unit 0 {
            family ethernet-switching {
                port-mode trunk;
                vlan {
                    members [10 20 30 40];
                }
            }
        }
    }
}
protocols {
    rstp {
        bridge-priority 16k;
            interface ge-0/0/13.0 {
                cost 1000;
                mode point-to-point;
            }
            interface ge-0/0/9.0 {
```

```
            cost 1000;
            mode point-to-point;
          }
          interface ge-0/0/11.0 {
            cost 1000;
            mode point-to-point;
          }
        }
      }
    }
    vlans {
      voice-vlan {
        vlan-id 10;
      }
      employee-vlan {
        vlan-id 20;
      }
      guest-vlan {
        vlan-id 30;
      }
      camera-vlan {
        vlan-id 40;
      }
    }
```

Configuring RSTP on Switch 2

To configure RSTP on switch 2, perform these tasks:

CLI Quick Configuration

To quickly configure interfaces and RSTP on Switch 2, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans voice-vlan description “Voice VLAN”
set vlans voice-vlan vlan-id 10
set vlans employee-vlan description “Employee VLAN”
set vlans employee-vlan vlan-id 20
set vlans guest-vlan description “Guest VLAN”
set vlans guest-vlan vlan-id 30
set vlans camera-vlan description “Camera VLAN”
set vlans camera-vlan vlan-id 40
set interfaces ge-0/0/14 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/18 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/14 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/18 unit 0 family ethernet-switching port-mode trunk
set protocols rstp bridge-priority 32k
set protocols rstp interface ge-0/0/14.0 cost 1000
set protocols rstp interface ge-0/0/14.0 mode point-to-point
set protocols rstp interface ge-0/0/18.0 cost 1000
set protocols rstp interface ge-0/0/18.0 mode point-to-point
```

Step-by-Step Procedure

To configure interfaces and RSTP on Switch 2:

1. Configure the VLANs **voice-vlan**, **employee-vlan**, **guest-vlan**, and **camera-vlan**:

```
[edit vlans]
user@switch2# set voice-vlan description “Voice VLAN”
user@switch2# set voice-vlan vlan-id 10
user@switch2# set employee-vlan description “Employee VLAN”
user@switch2# set employee-vlan vlan-id 20
user@switch2# set guest-vlan description “Guest VLAN”
user@switch2# set guest-vlan vlan-id 30
user@switch2# set camera-vlan vlan-description “Camera VLAN”
user@switch2# set guest-vlan vlan-id 40
```

2. Configure the VLANs on the interfaces, including support for the Ethernet Switching protocol:

```
[edit interfaces]
user@switch2# set ge-0/0/14 unit 0 family ethernet-switching vlan members [10 20
30 40]
user@switch2# set ge-0/0/18 unit 0 family ethernet-switching vlan members [10 20
30 40]
```

3. Configure the port mode for the interfaces:

```
[edit interfaces]
user@switch2# set ge-0/0/14 unit 0 family ethernet-switching port-mode trunk
user@switch2# set ge-0/0/18 unit 0 family ethernet-switching port-mode trunk
```

4. Configure RSTP on the switch:

```
[edit protocols]
user@switch2# rstp bridge-priority 32k
user@switch2# rstp interface ge-0/0/14.0 cost 1000
user@switch2# rstp interface ge-0/0/14.0 mode point-to-point
user@switch2# rstp interface ge-0/0/18.0 cost 1000
user@switch2# rstp interface ge-0/0/18.0 mode point-to-point
```

Results

Check the results of the configuration:

```
user@switch2> show configuration
interfaces {
    ge-0/0/14 {
        unit 0 {
            family ethernet-switching {
                port-mode trunk;
                vlan {
                    members [10 20 30 40];
                }
            }
        }
    }
    ge-0/0/18 {
        unit 0 {
            family ethernet-switching {
                port-mode trunk;
                vlan {
                    members [10 20 30 40];
                }
```

```
        }
      }
    }
  }
  protocols {
    rstp {
      bridge-priority 32k;
        interface ge-0/0/14.0 {
          cost 1000;
          mode point-to-point;
        }
        interface ge-0/0/18.0 {
          cost 1000;
          mode point-to-point;
        }
      }
    }
  }
  vlans {
    voice-vlan {
      vlan-id 10;
    }
    employee-vlan {
      vlan-id 20;
    }
    guest-vlan {
      vlan-id 30;
    }
    camera-vlan {
      vlan-id 40;
    }
  }
```

Configuring RSTP on Switch 3

To configure RSTP on switch 3, perform these tasks:

CLI Quick Configuration

To quickly configure interfaces and RSTP on Switch 3, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans voice-vlan description “Voice VLAN”
set vlans voice-vlan vlan-id 10
set vlans employee-vlan description “Employee VLAN”
set vlans employee-vlan vlan-id 20
set vlans guest-vlan description “Guest VLAN”
set vlans guest-vlan vlan-id 30
set vlans camera-vlan description “Camera VLAN”
set vlans camera-vlan vlan-id 40
set interfaces ge-0/0/26 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/28 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/24 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/26 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/28 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/24 unit 0 family ethernet-switching port-mode trunk
set protocols rstp bridge-priority 8k
set protocols rstp interface ge-0/0/26.0 cost 1000
```

```
set protocols rstp interface ge-0/0/26.0 mode point-to-point
set protocols rstp interface ge-0/0/28.0 cost 1000
set protocols rstp interface ge-0/0/28.0 mode point-to-point
set protocols rstp interface ge-0/0/24.0 cost 1000
set protocols rstp interface ge-0/0/24.0 mode point-to-point
```

Step-by-Step Procedure

To configure interfaces and RSTP on Switch 3:

1. Configure the VLANs **voice-vlan**, **employee-vlan**, **guest-vlan**, and **camera-vlan**:

```
[edit vlans]
user@switch3# set voice-vlan description “Voice VLAN”
user@switch3# set voice-vlan vlan-id 10
user@switch3# set employee-vlan description “Employee VLAN”
user@switch3# set employee-vlan vlan-id 20
user@switch3# set guest-vlan description “Guest VLAN”
user@switch3# set guest-vlan vlan-id 30
user@switch3# set camera-vlan description “Camera VLAN”
user@switch3# set guest-vlan vlan-id 40
```

2. Configure the VLANs on the interfaces, including support for the Ethernet Switching protocol:

```
[edit interfaces]
user@switch3# set ge-0/0/26 unit 0 family ethernet-switching vlan members [10 20
30 40]
user@switch3# set ge-0/0/28 unit 0 family ethernet-switching vlan members [10 20
30 40]
user@switch3# set ge-0/0/24 unit 0 family ethernet-switching vlan members [10 20
30 40]
```

3. Configure the port mode for the interfaces:

```
[edit interfaces]
user@switch3# set ge-0/0/26 unit 0 family ethernet-switching port-mode trunk
user@switch3# set ge-0/0/28 unit 0 family ethernet-switching port-mode trunk
user@switch3# set ge-0/0/24 unit 0 family ethernet-switching port-mode trunk
```

4. Configure RSTP on the switch:

```
[edit protocols]
user@switch3# rstp bridge-priority 8k
user@switch3# rstp interface ge-0/0/26.0 cost 1000
user@switch3# rstp interface ge-0/0/26.0 mode point-to-point
user@switch3# rstp interface ge-0/0/28.0 cost 1000
user@switch3# rstp interface ge-0/0/28.0 mode point-to-point
user@switch3# rstp interface ge-0/0/24.0 cost 1000
user@switch3# rstp interface ge-0/0/24.0 mode point-to-point
```

Results

Check the results of the configuration:

```
user@switch3> show configuration
  interfaces {
    ge-0/0/26 {
      unit 0 {
        family ethernet-switching {
          port-mode trunk;
          vlan {
            members [10 20 30 40];
          }
```

```
            }
          }
        }
        ge-0/0/28 {
          unit 0 {
            family ethernet-switching {
              port-mode trunk;
              vlan {
                members [10 20 30 40];
              }
            }
          }
        }
        ge-0/0/24 {
          unit 0 {
            family ethernet-switching {
              port-mode trunk;
              vlan {
                members [10 20 30 40];
              }
            }
          }
        }
      }
    }
    protocols {
      rstp {
        bridge-priority 8k;
          interface ge-0/0/26.0 {
            cost 1000;
            mode point-to-point;
          }
          interface ge-0/0/28.0 {
            cost 1000;
            mode point-to-point;
          }
          interface ge-0/0/24.0 {
            cost 1000;
            mode point-to-point;
          }
        }
        bridge-priority 8k;
          }
        }
      }
    }
    vlans {
      voice-vlan {
        vlan-id 10;
      }
      employee-vlan {
        vlan-id 20;
      }
      guest-vlan {
        vlan-id 30;
      }
```

```
    camera-vlan {
      vlan-id 40;
    }
  }
```

Configuring RSTP on Switch 4

To configure RSTP on switch 4, perform these tasks:

CLI Quick Configuration

To quickly configure interfaces and RSTP on Switch 4, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans voice-vlan description “Voice VLAN”
set vlans voice-vlan vlan–id 10
set vlans employee-vlan description “Employee VLAN”
set vlans employee-vlan vlan-id 20
set vlans guest-vlan description “Guest VLAN”
set vlans guest-vlan vlan-id 30
set vlans camera-vlan description “Camera VLAN”
set vlans camera-vlan vlan-id 40
set interfaces ge-0/0/23 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/19 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/23 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/19 unit 0 family ethernet-switching port-mode trunk
set protocols rstp bridge-priority 16k
set protocols rstp interface ge-0/0/23.0 cost 1000
set protocols rstp interface ge-0/0/23.0 mode point-to-point
set protocols rstp interface ge-0/0/19.0 cost 1000
set protocols rstp interface ge-0/0/19.0 mode point-to-point
```

Step-by-Step Procedure

To configure interfaces and RSTP on Switch 4:

1. Configure the VLANs **voice-vlan**, **employee-vlan**, **guest-vlan**, and **camera-vlan**:

   ```
   [edit vlans]
   user@switch4# set voice-vlan description “Voice VLAN”
   user@switch4# set voice-vlan vlan-id 10
   user@switch4# set employee-vlan description “Employee VLAN”
   user@switch4# set employee-vlan vlan-id 20
   user@switch4# set guest-vlan description “Guest VLAN”
   user@switch4# set guest-vlan vlan-id 30
   user@switch4# set camera-vlan description “Camera VLAN”
   user@switch4# set guest-vlan vlan-id 40
   ```

2. Configure the VLANs on the interfaces, including support for the Ethernet Switching protocol:

   ```
   [edit interfaces]
   user@switch4# set ge-0/0/23 unit 0 family ethernet-switching vlan members [10 20
   30 40]
   user@switch4# set ge-0/0/19 unit 0 family ethernet-switching vlan members [10 20
   30 40]
   ```

3. Configure the port mode for the interfaces:

   ```
   [edit interfaces]
   user@switch4# set ge-0/0/23 unit 0 family ethernet-switching port-mode trunk
   ```

```
user@switch4# set ge-0/0/19 unit 0 family ethernet-switching port-mode trunk
```

4. Configure RSTP on the switch:

```
[edit protocols]
user@switch4# rstp bridge-priority 16k
user@switch4# rstp interface all cost 1000
user@switch4# rstp interface ge-0/0/23.0 cost 1000
user@switch4# rstp interface ge-0/0/23.0 mode point-to-point
user@switch4# rstp interface ge-0/0/19.0 cost 1000
user@switch4# rstp interface ge-0/0/19.0 mode point-to-point
```

Results Check the results of the configuration:

```
user@switch4> show configuration
interfaces {
    ge-0/0/23 {
        unit 0 {
            family ethernet-switching {
                port-mode trunk;
                vlan {
                    members [10 20 30 40];
                }
            }
        }
    }
    ge-0/0/19 {
        unit 0 {
            family ethernet-switching {
                port-mode trunk;
                vlan {
                    members [10 20 30 40];
                }
            }
        }
    }
}
protocols {
    rstp {
        bridge-priority 16k;
            interface ge-0/0/23.0 {
                cost 1000;
                mode point-to-point;
            }
            interface ge-0/0/19.0 {
                cost 1000;
                mode point-to-point;
            }
        }
    }
}
vlans {
    voice-vlan {
        vlan-id 10;
    }
    employee-vlan {
```

```
    vlan-id 20;
  }
  guest-vlan {
    vlan-id 30;
  }
  camera-vlan {
    vlan-id 40;
  }
}
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying RSTP Configuration on Switch 1 on page 1295
- Verifying RSTP Configuration on Switch 2 on page 1295
- Verifying RSTP Configuration on Switch 3 on page 1296
- Verifying RSTP Configuration on Switch 4 on page 1296

Verifying RSTP Configuration on Switch 1

Purpose Verify the RSTP configuration on Switch 1.

Action Use the operational mode command:

```
user@switch1> show spanning-tree interface

Spanning tree interface parameters for instance 0

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ge-0/0/13.0    128:527      128:525  16384.0019e25040e0      1000  BLK    ALT
ge-0/0/9.0     128:529      128:513  32768.0019e2503d20      1000  BLK    ALT
ge-0/0/11.0    128:531      128:513   8192.0019e25051e0      1000  FWD    ROOT
```

Meaning Refer to the topology in Figure 33 on page 1284. The operational mode command **show spanning-tree interface** shows that **ge-0/0/13.0** is in a forwarding state. The other interfaces on Switch 1 are blocking.

Verifying RSTP Configuration on Switch 2

Purpose Verify the RSTP configuration on Switch 2.

Action Use the operational mode command:

```
user@switch2> show spanning-tree interface


Spanning tree interface parameters for instance 0

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ge-0/0/14.0     128:513      128:513  32768.0019e2503d20      1000  BLK    DESG
```

```
ge-0/0/18.0      128:519      128:515    8192.0019e25051e0      1000  FWD    ROOT
```

Meaning Refer to the topology in Figure 33 on page 1284. The operational mode command **show spanning-tree interface** shows that **ge-0/0/18.0** is in a forwarding state and the root port. The other interface on Switch 2 is blocking.

Verifying RSTP Configuration on Switch 3

Purpose Verify the RSTP configuration on Switch 3.

Action Use the operational mode commands:

```
user@switch3>  show spanning-tree interface

Spanning tree interface parameters for instance 0

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID         bridge ID          Cost
ge-0/0/26.0     128:513      128:513   8192.0019e25051e0      1000  FWD    DESG
ge-0/0/28.0     128:515      128:515   8192.0019e25051e0      1000  FWD    DESG
ge-0/0/24.0     128:517      128:517   8192.0019e25051e0      1000  FWD    DESG
```

Meaning Refer to the topology in Figure 33 on page 1284. The operational mode command **show spanning-tree interface** shows that no interface is the root interface.

Verifying RSTP Configuration on Switch 4

Purpose Verify the RSTP configuration on Switch 4.

Action Use the operational mode commands:

```
user@switch4>  show spanning-tree interface
Spanning tree interface parameters for instance 0

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID         bridge ID          Cost

ge-0/0/23.0    128:523      128:517   8192.0019e25051e0      1000  FWD    ROOT
ge-0/0/19.0    128:525      128:525  16384.0019e25040e0      1000  FWD    DESG
```

Meaning Refer to the topology in Figure 33 on page 1284. The operational mode command **show spanning-tree interface** shows that interface **ge-0/0/23.0** is the root interface and forwarding.

Related Documentation

- Example: Configuring Network Regions for VLANs with MSTP on J-EX Series Switches on page 1297
- Understanding RSTP for J-EX Series Switches on page 1276

Example: Configuring Network Regions for VLANs with MSTP on J-EX Series Switches

Multiple Spanning Tree Protocol (MSTP) is used to create a loop-free topology in networks using multiple spanning tree regions, each region containing multiple spanning-tree instances (MSTIs). MSTIs provide different paths for different VLANs. This functionality facilitates better load sharing across redundant links.

Up to 64 MSTI instances can be created for a J-EX Series switch, and each MSTI can support up to 4094 VLANs.

This example describes how to configure MSTP on four J-EX Series switches:

Requirements

This example uses the following hardware and software components:

- Four J-EX Series switches

Before you configure the switches for MSTP, be sure you have:

- Installed the four switches. See “Connecting and Configuring a J-EX Series Switch (J-Web Procedure)” on page 163.
- Performed the initial software configuration on all switches. See Installing and Connecting a J-EX4200 Switch.

Overview and Topology

When the number of VLANs grows in a network, MSTP provides a more efficient way of creating a loop-free topology using MSTIs. Each MSTI in the spanning tree domain maintains its own tree. Each tree can be mapped to different links, utilizing bandwidth that would be unavailable to a single tree. MSTIs reduce demand on system resources.

Figure 34: Network Topology for MSTP

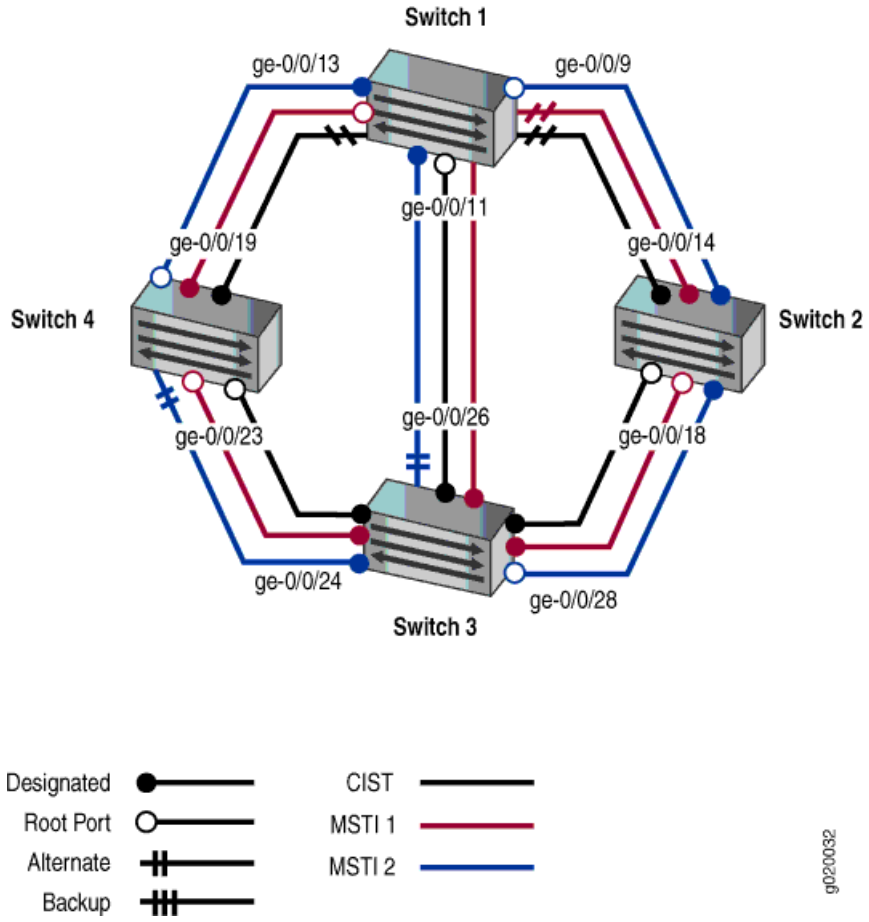

The interfaces shown in Table 172 on page 1298 will be configured for MSTP.

NOTE: You can configure MSTP on logical or physical interfaces. This example shows MSTP configured on logical interfaces.

Table 172: Components of the Topology for Configuring MSTP on J-EX Series Switches

Property	Settings
Switch 1	The following ports on Switch 1 are connected in this way: • **ge-0/0/9** is connected to Switch 2 • **ge-0/0/13** is connected to Switch 4 • **ge-0/0/11** is connected to Switch 3
Switch 2	The following ports on Switch 2 are connected in this way: • **ge-0/0/14** is connected to Switch 1 • **ge-0/0/18** is connected to Switch 3

Table 172: Components of the Topology for Configuring MSTP on J-EX Series Switches *(continued)*

Property	Settings
Switch 3	The following ports on Switch 3 are connected in this way: • **ge-0/0/26** is connected to Switch 1 • **ge-0/0/28** is connected to Switch 2 • **ge-0/0/24** is connected to Switch 4
Switch 4	The following ports on Switch 4 are connected in this way: • **ge-0/0/19** is connected to Switch 1 • **ge-0/0/23** is connected to Switch 3
VLAN names and tag IDs	**voice-vlan**, tag **10** **employee-vlan**, tag **20** **guest-vlan**, tag **30** **camera-vlan**, tag **40**
MSTIs	**1** **2**

The topology in Figure 34 on page 1298 shows a Common Internal Spanning Tree (CIST). The CIST is a single spanning tree connecting all devices in the network. The switch with the highest priority is elected as the root bridge of the CIST.

Also in an MSTP topology are ports that have specific roles:

- The root port is responsible for forwarding data to the root bridge.
- The alternate port is a standby port for the root port. When a root port goes down, the alternate port becomes the active root port.
- The designated port forwards data to the downstream network segment or device.
- The backup port is a backup port for the designated port. When a designated port goes down, the backup port becomes the active designated port and starts forwarding data.

In this example, one MSTP region, **region1**, contains Switch 1, Switch 2, Switch 3, and Switch 4. Within the region, four VLANs are created:

- The **voice-vlan** supports voice traffic and has a VLAN tag identifier of 10.
- **employee-vlan** supports data traffic and has a VLAN tag identifier of 20.
- The **guest-vlan** supports guest VLAN traffic (for supplicants that fail 802-1X authentication) and has a VLAN tag identifier of 30.
- The **camera-vlan** supports video traffic and has a VLAN tag identifier of 40.

The VLANs are associated with specific interfaces on each of the four switches. Two MSTIs, **1** and **2**, are then associated with the VLAN tag identifiers, and some MSTP parameters, such as cost, are configured on each switch.

Configuring MSTP on Switch 1

To configure MSTP on Switch 1, perform these tasks:

CLI Quick Configuration

To quickly configure interfaces and MSTP on Switch 1, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans voice-vlan description “Voice VLAN”
set vlans voice-vlan vlan-id 10
set vlans employee-vlan description “Employee VLAN”
set vlans employee-vlan vlan-id 20
set vlans guest-vlan description “Guest VLAN”
set vlans guest-vlan vlan-id 30
set vlans camera-vlan description “Camera VLAN”
set vlans camera-vlan vlan-id 40
set interfaces ge-0/0/13 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge–0/0/9 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge–0/0/11 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge–0/0/13 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/9 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/11 unit 0 family ethernet-switching port-mode trunk
set protocols mstp configuration-name region1
set protocols mstp bridge-priority 16k
set protocols mstp interface ge-0/0/13.0 cost 1000
set protocols mstp interface ge-0/0/13.0 mode point-to-point
set protocols mstp interface ge-0/0/9.0 cost 1000
set protocols mstp interface ge-0/0/9.0 mode point-to-point
set protocols mstp interface ge-0/0/11.0 cost 1000
set protocols mstp interface ge-0/0/11.0 mode point-to-point
set protocols mstp msti 1 bridge-priority 16k
set protocols mstp msti 1 vlan [10 20]
set protocols mstp msti 1 interface ge-0/0/11.0 cost 4000
set protocols mstp msti 2 bridge-priority 8k
set protocols mstp msti 2 vlan [30 40]
```

Step-by-Step Procedure

To configure interfaces and MSTP on Switch 1:

1. Configure the VLANs **voice-vlan**, **employee-vlan**, **guest-vlan**, and **camera-vlan**:

   ```
   [edit vlans]
   user@switch1# set voice-vlan description “Voice VLAN”
   user@switch1# set voice-vlan vlan-id 10
   user@switch1# set employee-vlan description “Employee VLAN”
   user@switch1# set employee-vlan vlan-id 20
   user@switch1# set guest-vlan description “Guest VLAN”
   user@switch1# set guest-vlan vlan-id 30
   user@switch1# set camera-vlan description “Camera VLAN”
   user@switch1# set guest-vlan vlan-id 40
   ```

2. Configure the VLANs on the interfaces, including support for the Ethernet Switching protocol:

   ```
   [edit interfaces]
   user@switch1# set ge–0/0/13 unit 0 family ethernet-switching vlan members [10 20
   30 40]
   user@switch1# set ge-0/0/9 unit 0 family ethernet-switching vlan members [10 20 30
   40]
   user@switch1# set ge-0/0/11 unit 0 family ethernet-switching vlan members [10 20
   30 40]
   ```

3. Configure the port mode for the interfaces:

```
[edit interfaces]
user@switch1# set ge–0/0/13 unit 0 family ethernet-switching port-mode trunk
user@switch1# set ge-0/0/9 unit 0 family ethernet-switching port-mode trunk
user@switch1# set ge-0/0/11 unit 0 family ethernet-switching port-mode trunk
```

4. Configure MSTP on the switch, including the two MSTIs:

```
[edit protocols]
user@switch1# mstp configuration-name region1
user@switch1# mstp bridge-priority 16k
user@switch1# mstp interface ge-0/0/13.0 cost 1000
user@switch1# mstp interface ge-0/0/13.0 mode point-to-point
user@switch1# mstp interface ge-0/0/9.0 cost 1000
user@switch1# mstp interface ge-0/0/9.0 mode point-to-point
user@switch1# mstp interface ge-0/0/11.0 cost 4000
user@switch1# mstp interface ge-0/0/11.0 mode point-to-point
user@switch1# mstp msti 1 bridge-priority 16k
user@switch1# mstp msti 1 vlan [10 20]
user@switch1# mstp msti 1 interface ge-0/0/11.0 cost 4000
user@switch1# mstp msti 2 bridge-priority 8k
user@switch1# mstp msti 2 vlan [30 40]
```

Results Check the results of the configuration:

```
user@switch1> show configuration
interfaces {
  ge-0/0/13 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members 10;
          members 20;
          members 30;
          members 40;
        }
      }
    }
  }
  ge-0/0/9 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members 10;
          members 20;
          members 30;
          members 40;
        }
      }
    }
  }
  ge-0/0/11 {
    unit 0 {
      family ethernet-switching {
```

```
            port-mode trunk;
            vlan {
              members 10;
              members 20;
              members 30;
              members 40;
            }
          }
        }
      }
    }
    protocols {
      mstp {
        configuration-name region1;
      bridge-priority 16k;
        interface ge-0/0/13.0 {
          cost 1000;
          mode point-to-point;
        }
        interface ge-0/0/9.0 {
          cost 1000;
          mode point-to-point;
        }
        interface ge-0/0/11.0 {
          cost 4000;
          mode point-to-point;
        }
      msti 1 {
        bridge-priority 16k;
        vlan [ 10 20 ];
        interface ge-0/0/11.0 {
          cost 4000;
        }
      }
      msti 2 {
        bridge-priority 8k;
        vlan [ 30 40 ];
      }
    }
    vlans {
      voice-vlan {
        vlan-id 10;
      }
      employee-vlan {
        vlan-id 20;
      }
      guest-vlan {
        vlan-id 30;
      }
      camera-vlan {
        vlan-id 40;
      }
    }
```

Configuring MSTP on Switch 2

To configure on Switch 2, perform these tasks:

CLI Quick Configuration

To quickly configure interfaces and MSTP on Switch 2, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans voice-vlan description “Voice VLAN”
set vlans voice-vlan vlan—id 10
set vlans employee-vlan description “Employee VLAN”
set vlans employee-vlan vlan-id 20
set vlans guest-vlan description “Guest VLAN”
set vlans guest-vlan vlan-id 30
set vlans camera-vlan description “Camera VLAN”
set vlans camera-vlan vlan-id 40
set interfaces ge—0/0/14 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/18 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/14 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/18 unit 0 family ethernet-switching port-mode trunk
set protocols mstp configuration-name region1
set protocols mstp bridge-priority 32k
set protocols mstp interface ge-0/0/14.0 cost 1000
set protocols mstp interface ge-0/0/14.0 mode point-to-point
set protocols mstp interface ge-0/0/18.0 cost 1000
set protocols mstp interface ge-0/0/18.0 mode point-to-point
set protocols mstp msti 1 bridge-priority 32k
set protocols mstp msti 1 vlan [10 20]
set protocols mstp msti 2 bridge-priority 4k
set protocols mstp msti 2 vlan [30 40]
```

Step-by-Step Procedure

To configure interfaces and MSTP on Switch 2:

1. Configure the VLANs **voice-vlan**, **employee-vlan**, **guest-vlan**, and **camera-vlan**:

   ```
   [edit vlans]
   user@switch2# set voice-vlan description “Voice VLAN”
   user@switch2# set voice-vlan vlan-id 10
   user@switch2# set employee-vlan description “Employee VLAN”
   user@switch2# set employee-vlan vlan-id 20
   user@switch2# set guest-vlan description “Guest VLAN”
   user@switch2# set guest-vlan vlan-id 30
   user@switch2# set camera-vlan vlan-description “Camera VLAN”
   user@switch2# set guest-vlan vlan-id 40
   ```

2. Configure the VLANs on the interfaces, including support for the Ethernet Switching protocol:

   ```
   [edit interfaces]
   user@switch2# set ge-0/0/14 unit 0 family ethernet-switching vlan members [10 20
   30 40]
   user@switch2# set ge-0/0/18 unit 0 family ethernet-switching vlan members [10 20
   30 40]
   ```

3. Configure the port mode for the interfaces:

   ```
   [edit interfaces]
   user@switch2# set ge-0/0/14 unit 0 family ethernet-switching port-mode trunk
   user@switch2# set ge-0/0/18 unit 0 family ethernet-switching port-mode trunk
   ```

4. Configure MSTP on the switch, including the two MSTIs:

```
[edit protocols]
user@switch2# mstp configuration-name region1
user@switch2# mstp bridge-priority 32k
user@switch2# mstp interface ge-0/0/14.0 cost 1000
user@switch2# mstp interface ge-0/0/14.0 mode point-to-point
user@switch2# mstp interface ge-0/0/18.0 cost 1000
user@switch2# mstp interface ge-0/0/18.0 mode point-to-point
user@switch2# mstp interface all cost 1000
user@switch2# mstp msti 1 bridge-priority 32k
user@switch2# mstp msti 1 vlan [10 20]
user@switch2# mstp msti 2 bridge-priority 4k
user@switch2# mstp msti 2 vlan [30 40]
```

Results Check the results of the configuration:

```
user@switch2> show configuration
interfaces {
  ge-0/0/14 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members 10;
          members 20;
          members 30;
          members 40;
        }
      }
    }
  }
  ge-0/0/18 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members 10;
          members 20;
          members 30;
          members 40;
        }
      }
    }
  }
}
protocols {
  mstp {
    configuration-name region1;
    bridge-priority 32k;
    interface ge-0/0/14.0 {
      cost 1000;
      mode point-to-point;
    }
    interface ge-0/0/18.0 {
      cost 1000;
```

```
        mode point-to-point;
      }
      msti 1 {
        bridge-priority 32k;
        vlan [ 10 20 ];
      }
      msti 2 {
        bridge-priority 4k;
        vlan [ 30 40 ];
      }
    }
  }
  vlans {
    voice-vlan {
      vlan-id 10;
    }
    employee-vlan {
      vlan-id 20;
    }
    guest-vlan {
      vlan-id 30;
    }
    camera-vlan {
      vlan-id 40;
    }
  }
```

Configuring MSTP on Switch 3

To configure MSTP on Switch 3, perform these tasks:

CLI Quick Configuration

To quickly configure interfaces and MSTP on Switch 3, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans voice-vlan description “Voice VLAN”
set vlans voice-vlan vlan-id 10
set vlans employee-vlan description “Employee VLAN”
set vlans employee-vlan vlan-id 20
set vlans guest-vlan description “Guest VLAN”
set vlans guest-vlan vlan-id 30
set vlans camera-vlan description “Camera VLAN”
set vlans camera-vlan vlan—id 40
set interfaces ge—0/0/26 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge—0/0/28 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge—0/0/24 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/26 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/28 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/24 unit 0 family ethernet-switching port-mode trunk
set protocols mstp configuration-name region1
set protocols mstp bridge-priority 8k
set protocols mstp interface ge-0/0/26.0 cost 1000
set protocols mstp interface ge-0/0/26.0 mode point-to-point
set protocols mstp interface ge-0/0/28.0 cost 1000
set protocols mstp interface ge-0/0/28.0 mode point-to-point
set protocols mstp interface ge-0/0/24.0 cost 1000
set protocols mstp interface ge-0/0/24.0 mode point-to-point
```

```
set protocols mstp msti 1 bridge-priority 4k
set protocols mstp msti 1 vlan [10 20]
set protocols mstp msti 2 bridge-priority 16k
set protocols mstp msti 2 vlan [30 40]
```

Step-by-Step Procedure

To configure interfaces and MSTP on Switch 3:

1. Configure the VLANs **voice-vlan**, **employee-vlan**, **guest-vlan**, and **camera-vlan**:

```
[edit vlans]
user@switch3# set voice-vlan description “Voice VLAN”
user@switch3# set voice-vlan vlan-id 10
user@switch3# set employee-vlan description “Employee VLAN”
user@switch3# set employee-vlan vlan-id 20
user@switch3# set guest-vlan description “Guest VLAN”
user@switch3# set guest-vlan vlan-id 30
user@switch3# set camera-vlan description “Camera VLAN”
user@switch3# set guest-vlan vlan-id 40
```

2. Configure the VLANs on the interfaces, including support for the Ethernet Switching protocol:

```
[edit interfaces]
user@switch3# set ge-0/0/26 unit 0 family ethernet-switching vlan members [10 20
30 40]
user@switch3# set ge-0/0/28 unit 0 family ethernet-switching vlan members [10 20
30 40]
user@switch3# set ge-0/0/24 unit 0 family ethernet-switching vlan members [10 20
30 40]
```

3. Configure the port mode for the interfaces:

```
[edit interfaces]
user@switch3# set ge-0/0/26 unit 0 family ethernet-switching port-mode trunk
user@switch3# set ge-0/0/28 unit 0 family ethernet-switching port-mode trunk
user@switch3# set ge-0/0/24 unit 0 family ethernet-switching port-mode trunk
```

4. Configure MSTP on the switch, including the two MSTIs:

```
[edit protocols]
user@switch3# mstp configuration-name region1
user@switch3# mstp bridge-priority 8k
user@switch3# mstp interface ge-0/0/26.0 cost 1000
user@switch3# mstp interface ge-0/0/26.0 mode point-to-point
user@switch3# mstp interface ge-0/0/28.0 cost 1000
user@switch3# mstp interface ge-0/0/28.0 mode point-to-point
user@switch3# mstp interface ge-0/0/24.0 cost 1000
user@switch3# mstp interface ge-0/0/24.0 mode point-to-point
user@switch3# mstp interface all cost 1000
user@switch3# mstp msti 1 bridge-priority 4k
user@switch3# mstp msti 1 vlan [10 20]
user@switch3# mstp msti 2 bridge-priority 16k
user@switch3# mstp msti 2 vlan [30 40]
```

Results

Check the results of the configuration:

```
user@switch3> show configuration
  interfaces {
    ge-0/0/26 {
      unit 0 {
```

```
            family ethernet-switching {
              port-mode trunk;
              vlan {
                members 10;
                members 20;
                members 30;
                members 40;
              }
            }
          }
        }
        ge-0/0/28 {
          unit 0 {
            family ethernet-switching {
              port-mode trunk;
              vlan {
                members 10;
                members 20;
                members 30;
                members 40;
              }
            }
          }
        }
        ge-0/0/24 {
          unit 0 {
            family ethernet-switching {
              port-mode trunk;
              vlan {
                members 10;
                members 20;
                members 30;
                members 40;
              }
            }
          }
        }
      }
    }
    protocols {
      mstp {
        configuration-name region1;
        bridge-priority 8k;
        interface ge-0/0/26.0 {
          cost 1000;
          mode point-to-point;
        }
        interface ge-0/0/28.0 {
          cost 1000;
          mode point-to-point;
        }
        interface ge-0/0/24.0 {
          cost 1000;
          mode point-to-point;
        }
        msti 1 {
```

```
        bridge-priority 4k;
        vlan [ 10 20 ];
      }
      msti 2 {
        bridge-priority 16k;
        vlan [ 30 40 ];
      }
    }
  }
  vlans {
    voice-vlan {
      vlan-id 10;
    }
    employee-vlan {
      vlan-id 20;
    }
    guest-vlan {
      vlan-id 30;
    }
    camera-vlan {
      vlan-id 40;
    }
  }
```

Configuring MSTP on Switch 4

To configure MSTP on Switch 4, perform these tasks:

CLI Quick Configuration

To quickly configure interfaces and MSTP on Switch 4, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans voice–vlan description “Voice VLAN”
set vlans voice-vlan vlan–id 10
set vlans employee—vlan description “Employee VLAN”
set vlans employee—vlan vlan—id 20
set vlans guest—vlan description “Guest VLAN”
set vlans guest—vlan vlan—id 30
set vlans camera—vlan description “Camera VLAN”
set vlans camera—vlan vlan—id 40
set interfaces ge-0/0/23 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge-0/0/19 unit 0 family ethernet-switching vlan members [10 20 30 40]
set interfaces ge—0/0/23 unit 0 family ethernet-switching port-mode trunk
set interfaces ge—0/0/19 unit 0 family ethernet-switching port-mode trunk
set protocols mstp configuration-name region1
set protocols mstp bridge-priority 16k
set protocols mstp interface ge—0/0/23.0 cost 1000
set protocols mstp interface ge—0/0/23.0 mode point-to-point
set protocols mstp interface ge—0/0/19.0 cost 1000
set protocols mstp interface ge—0/0/19.0 mode point-to-point
set protocols mstp msti 1 bridge-priority 16k
set protocols mstp msti 1 vlan [10 20]
set protocols mstp msti 2 bridge-priority 32k
set protocols mstp msti 2 vlan [30 40]
```

Step-by-Step Procedure

To configure interfaces and MSTP on Switch 4:

1. Configure the VLANs **voice-vlan**, **employee-vlan**, **guest-vlan**, and **camera-vlan**:

```
[edit vlans]
user@switch4# set voice-vlan description “Voice VLAN”
user@switch4# set voice-vlan vlan—id 10
user@switch4# set employee-vlan description “Employee VLAN”
user@switch4# set employee-vlan vlan—id 20
user@switch4# set guest-vlan description “Guest VLAN”
user@switch4# set guest-vlan vlan—id 30
user@switch4# set camera-vlan description “Camera VLAN”
user@switch4# set guest-vlan vlan—id 40
```

2. Configure the VLANs on the interfaces, including support for the Ethernet Switching protocol:

```
[edit interfaces]
user@switch4# set ge-0/0/23 unit 0 family ethernet-switching vlan members [10 20
30 40]
user@switch4# set ge-0/0/19 unit 0 family ethernet-switching vlan members [10 20
30 40]
```

3. Configure the port mode for the interfaces:

```
[edit interfaces]
user@switch4# set ge—0/0/23 unit 0 family ethernet-switching port-mode trunk
user@switch4# set ge—0/0/19 unit 0 family ethernet-switching port-mode trunk
```

4. Configure MSTP on the switch, including the two MSTIs:

```
[edit protocols]
user@switch4# mstp configuration-name region1
user@switch4# mstp bridge-priority 16k
user@switch4# mstp interface all cost 1000
user@switch4# mstp interface ge—0/0/23.0 cost 1000
user@switch4# mstp interface ge—0/0/23.0 mode point-to-point
user@switch4# mstp interface ge—0/0/19.0 cost 1000
user@switch4# mstp interface ge—0/0/19.0 mode point-to-point
user@switch4# mstp msti 1 bridge-priority 16k
user@switch4# mstp msti 1 vlan [10 20]
user@switch4# mstp msti 2 bridge-priority 32k
user@switch4# mstp msti 2 vlan [30 40]
```

Results

Check the results of the configuration:

```
user@switch4> show configuration
interfaces {
  ge-0/0/23 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members 10;
          members 20;
          members 30;
          members 40;
        }
      }
```

```
        }
    }
    ge-0/0/19 {
        unit 0 {
            family ethernet-switching {
                port-mode trunk;
                vlan {
                    members 10;
                    members 20;
                    members 30;
                    members 40;
                }
            }
        }
    }
}
protocols {
    mstp {
        configuration-name region1;
        bridge-priority 16k;
        interface ge-0/0/23.0 {
            cost 1000;
            mode point-to-point;
        }
        interface ge-0/0/19.0 {
            cost 1000;
            mode point-to-point;
        }
        msti 1 {
            bridge-priority 16k;
            vlan [ 10 20 ];
        }
        msti 2 {
            bridge-priority 32k;
            vlan [ 30 40 ];
        }
    }
}
vlans {
    voice-vlan {
        vlan-id 10;
    }
    employee-vlan {
        vlan-id 20;
    }
    guest-vlan {
        vlan-id 30;
    }
    camera-vlan {
        vlan-id 40;
    }
}
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying MSTP Configuration on Switch 1 on page 1311
- Verifying MSTP Configuration on Switch 2 on page 1312
- Verifying MSTP Configuration on Switch 3 on page 1314
- Verifying MSTP Configuration on Switch 4 on page 1315

Verifying MSTP Configuration on Switch 1

Purpose Verify the MSTP configuration on Switch 1.

Action Use the operational mode commands:

```
user@switch1> show spanning-tree interface
Spanning tree interface parameters for instance 0

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ge-0/0/13.0    128:527      128:525  16384.0019e25040e0      1000  FWD    ROOT
ge-0/0/9.0     128:529      128:513  32768.0019e2503d20      1000  BLK    ALT
ge-0/0/11.0    128:531      128:513   8192.0019e25051e0      4000  BLK    ALT

Spanning tree interface parameters for instance 1

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ge-0/0/13.0    128:527      128:525  16385.0019e25040e0      1000  FWD    ROOT
ge-0/0/9.0     128:529      128:513  32769.0019e2503d20      1000  BLK    ALT
ge-0/0/11.0    128:531      128:513   4097.0019e25051e0      4000  BLK    ALT

Spanning tree interface parameters for instance 2

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ge-0/0/13.0    128:527      128:527   8194.0019e25044e0      1000  FWD    DESG
ge-0/0/9.0     128:529      128:513   4098.0019e2503d20      1000  FWD    ROOT
ge-0/0/11.0    128:531      128:531   8194.0019e25044e0      1000  FWD    DESG


user@switch1> show spanning-tree bridge
STP bridge parameters
Context ID                          : 0
Enabled protocol                    : MSTP

STP bridge parameters for CIST
  Root ID                           : 8192.00:19:e2:50:51:e0
  Root cost                         : 0
  Root port                         : ge-0/0/13.0
  CIST regional root                : 8192.00:19:e2:50:51:e0
  CIST internal root cost           : 2000
  Hello time                        : 2 seconds
  Maximum age                       : 20 seconds
  Forward delay                     : 15 seconds
  Hop count                         : 18
  Message age                       : 0
  Number of topology changes        : 3
```

```
  Time since last topology change   : 921 seconds
  Local parameters
    Bridge ID                       : 16384.00:19:e2:50:44:e0
    Extended system ID              : 0
    Internal instance ID            : 0

STP bridge parameters for MSTI 1
  MSTI regional root                : 4097.00:19:e2:50:51:e0
  Root cost                         : 2000
  Root port                         : ge-0/0/13.0
  Hello time                        : 2 seconds
  Maximum age                       : 20 seconds
  Forward delay                     : 15 seconds
  Hop count                         : 18
  Local parameters
    Bridge ID                       : 16385.00:19:e2:50:44:e0
    Extended system ID              : 0
    Internal instance ID            : 1

STP bridge parameters for MSTI 2
  MSTI regional root                : 4098.00:19:e2:50:3d:20
  Root cost                         : 1000
  Root port                         : ge-0/0/9.0
  Hello time                        : 2 seconds
  Maximum age                       : 20 seconds
  Forward delay                     : 15 seconds
  Hop count                         : 19
  Local parameters
    Bridge ID                       : 8194.00:19:e2:50:44:e0
    Extended system ID              : 0
    Internal instance ID            : 2
```

Meaning The operational mode command **show spanning-tree interface** displays spanning-tree domain information such as the designated port and the port roles.

The operational mode command **show spanning-tree bridge** displays the spanning-tree domain information at either the bridge level or interface level. If the optional interface name is omitted, all interfaces in the spanning-tree domain are displayed.

Verifying MSTP Configuration on Switch 2

Purpose Verify the MSTP configuration on Switch 2.

Action Use the operational mode commands:

```
user@switch2>  show spanning-tree interface

Spanning tree interface parameters for instance 0

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ge-0/0/14.0     128:513      128:513  32768.0019e2503d20      1000  FWD    DESG
ge-0/0/18.0     128:519      128:515   8192.0019e25051e0      1000  FWD    ROOT

Spanning tree interface parameters for instance 1

Interface    Port ID    Designated      Designated         Port    State  Role
```

```
                          port ID        bridge ID          Cost
ge-0/0/14.0     128:513      128:513  32769.0019e2503d20      1000  FWD    DESG
ge-0/0/18.0     128:519      128:515   4097.0019e25051e0      1000  FWD    ROOT

Spanning tree interface parameters for instance 2

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ge-0/0/14.0     128:513      128:513   4098.0019e2503d20      1000  FWD    DESG

ge-0/0/18.0     128:519      128:519   4098.0019e2503d20      1000  FWD    DESG


user@switch2> show spanning-tree bridge
STP bridge parameters
Context ID                          : 0
Enabled protocol                    : MSTP

STP bridge parameters for CIST
  Root ID                           : 8192.00:19:e2:50:51:e0
  Root cost                         : 0
  Root port                         : ge-0/0/18.0
  CIST regional root                : 8192.00:19:e2:50:51:e0
  CIST internal root cost           : 1000
  Hello time                        : 2 seconds
  Maximum age                       : 20 seconds
  Forward delay                     : 15 seconds
  Hop count                         : 19
  Message age                       : 0
  Number of topology changes        : 1
  Time since last topology change   : 782 seconds
  Local parameters
    Bridge ID                       : 32768.00:19:e2:50:3d:20
    Extended system ID              : 0
    Internal instance ID            : 0

STP bridge parameters for MSTI 1
  MSTI regional root                : 4097.00:19:e2:50:51:e0
  Root cost                         : 1000
  Root port                         : ge-0/0/18.0
  Hello time                        : 2 seconds
  Maximum age                       : 20 seconds
  Forward delay                     : 15 seconds
  Hop count                         : 19
  Local parameters
    Bridge ID                       : 32769.00:19:e2:50:3d:20
    Extended system ID              : 0
    Internal instance ID            : 1

STP bridge parameters for MSTI 2
  MSTI regional root                : 4098.00:19:e2:50:3d:20
  Hello time                        : 2 seconds
  Maximum age                       : 20 seconds
  Forward delay                     : 15 seconds
  Local parameters
    Bridge ID                       : 4098.00:19:e2:50:3d:20
    Extended system ID              : 0
    Internal instance ID            : 2
```

Meaning The operational mode command **show spanning-tree interface** displays spanning-tree domain information such as the designated port and the port roles.

The operational mode command **show spanning-tree bridge** displays the spanning-tree domain information at either the bridge level or interface level. If the optional interface name is omitted, all interfaces in the spanning-tree domain are displayed.

Verifying MSTP Configuration on Switch 3

Purpose Verify the MSTP configuration on Switch 3.

Action Use the operational mode commands:

```
user@switch3>  show spanning-tree interface

Spanning tree interface parameters for instance 0

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ge-0/0/26.0     128:513      128:513   8192.0019e25051e0      1000  FWD    DESG
ge-0/0/28.0     128:515      128:515   8192.0019e25051e0      1000  FWD    DESG
ge-0/0/24.0     128:517      128:517   8192.0019e25051e0      1000  FWD    DESG

Spanning tree interface parameters for instance 1

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ge-0/0/26.0     128:513      128:513   4097.0019e25051e0      1000  FWD    DESG
ge-0/0/28.0     128:515      128:515   4097.0019e25051e0      1000  FWD    DESG
ge-0/0/24.0     128:517      128:517   4097.0019e25051e0      1000  FWD    DESG


Spanning tree interface parameters for instance 2

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ge-0/0/26.0     128:513      128:531   8194.0019e25044e0      1000  BLK    ALT
ge-0/0/28.0     128:515      128:519   4098.0019e2503d20      1000  FWD    ROOT
ge-0/0/24.0     128:517      128:517  16386.0019e25051e0      1000  FWD    DESG


user@switch3>  show spanning-tree bridge
STP bridge parameters
Context ID                          : 0
Enabled protocol                    : MSTP

STP bridge parameters for CIST
  Root ID                           : 8192.00:19:e2:50:51:e0
  CIST regional root                : 8192.00:19:e2:50:51:e0
  CIST internal root cost           : 0
  Hello time                        : 2 seconds
  Maximum age                       : 20 seconds
  Forward delay                     : 15 seconds
  Number of topology changes        : 3
  Time since last topology change   : 843 seconds
  Local parameters
    Bridge ID                       : 8192.00:19:e2:50:51:e0
    Extended system ID              : 0
    Internal instance ID            : 0

STP bridge parameters for MSTI 1
  MSTI regional root                : 4097.00:19:e2:50:51:e0
  Hello time                        : 2 seconds
```

```
  Maximum age                        : 20 seconds
  Forward delay                      : 15 seconds
  Local parameters
    Bridge ID                        : 4097.00:19:e2:50:51:e0
    Extended system ID               : 0
    Internal instance ID             : 1

STP bridge parameters for MSTI 2
  MSTI regional root                 : 4098.00:19:e2:50:3d:20
  Root cost                          : 1000
  Root port                          : ge-0/0/28.0
  Hello time                         : 2 seconds
  Maximum age                        : 20 seconds
  Forward delay                      : 15 seconds
  Hop count                          : 19
  Local parameters
    Bridge ID                        : 16386.00:19:e2:50:51:e0
    Extended system ID               : 0
    Internal instance ID             : 2
```

Meaning The operational mode command **show spanning-tree interface** displays spanning-tree domain information such as the designated port and the port roles.

The operational mode command **show spanning-tree bridge** displays the spanning-tree domain information at either the bridge level or interface level. If the optional interface name is omitted, all interfaces in the spanning-tree domain are displayed.

Verifying MSTP Configuration on Switch 4

Purpose Verify the MSTP configuration on Switch 4.

Action Use the operational mode commands:

```
user@switch4> show spanning-tree interface
Spanning tree interface parameters for instance 0

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost

ge-0/0/23.0    128:523      128:517   8192.0019e25051e0      1000  FWD    ROOT
ge-0/0/19.0    128:525      128:525  16384.0019e25040e0      1000  FWD    DESG

Spanning tree interface parameters for instance 1

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ge-0/0/23.0    128:523      128:517   4097.0019e25051e0      1000  FWD    ROOT
ge-0/0/19.0    128:525      128:525  16385.0019e25040e0      1000  FWD    DESG


Spanning tree interface parameters for instance 2

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ge-0/0/23.0    128:523      128:517  16386.0019e25051e0      1000  BLK    ALT
ge-0/0/19.0    128:525      128:527   8194.0019e25044e0      1000  FWD    ROOT


user@switch4> show spanning-tree bridge
```

```
STP bridge parameters
Context ID                          : 0
Enabled protocol                    : MSTP

STP bridge parameters for CIST
  Root ID                           : 8192.00:19:e2:50:51:e0
  Root cost                         : 0
  Root port                         : ge-0/0/23.0
  CIST regional root                : 8192.00:19:e2:50:51:e0
  CIST internal root cost           : 1000
  Hello time                        : 2 seconds
  Maximum age                       : 20 seconds
  Forward delay                     : 15 seconds
  Hop count                         : 19
  Message age                       : 0
  Number of topology changes        : 4
  Time since last topology change   : 887 seconds
  Local parameters
    Bridge ID                       : 16384.00:19:e2:50:40:e0
    Extended system ID              : 0
    Internal instance ID            : 0

STP bridge parameters for MSTI 1
  MSTI regional root                : 4097.00:19:e2:50:51:e0
  Root cost                         : 1000
  Root port                         : ge-0/0/23.0
  Hello time                        : 2 seconds
  Maximum age                       : 20 seconds
  Forward delay                     : 15 seconds
  Hop count                         : 19
  Local parameters
    Bridge ID                       : 16385.00:19:e2:50:40:e0
    Extended system ID              : 0
    Internal instance ID            : 1

STP bridge parameters for MSTI 2
  MSTI regional root                : 4098.00:19:e2:50:3d:20
  Root cost                         : 2000
  Root port                         : ge-0/0/19.0
  Hello time                        : 2 seconds
  Maximum age                       : 20 seconds
  Forward delay                     : 15 seconds
  Hop count                         : 18
  Local parameters
    Bridge ID                       : 32770.00:19:e2:50:40:e0
    Extended system ID              : 0
    Internal instance ID            : 2
```

Meaning The operational mode command **show spanning-tree interface** displays spanning-tree domain information such as the designated port and the port roles.

The operational mode command **show spanning-tree bridge** displays the spanning-tree domain information at either the bridge level or interface level. If the optional interface name is omitted, all interfaces in the spanning-tree domain are displayed.

Related Documentation

- Example: Configuring Faster Convergence and Improving Network Stability with RSTP on J-EX Series Switches on page 1283
- Understanding MSTP for J-EX Series Switches on page 1277

Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations on J-EX Series Switches

J-EX Series switches provide Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP). Configure BPDU protection on interfaces to prevent them from receiving BPDUs that could result in STP misconfigurations, which could lead to network outages.

This example describes how to configure BPDU protection on access interfaces on a J-EX Series switch in an RSTP topology:

Requirements

This example uses the following hardware and software components:

- Two J-EX Series switches in an RSTP topology

Before you configure the interfaces on Switch 2 for BPDU protection, be sure you have:

- RSTP operating on the switches.

NOTE: By default, RSTP is enabled on all J-EX Series switches.

Overview and Topology

A loop-free network is supported through the exchange of a special type of frame called bridge protocol data unit (BPDU). Receipt of BPDUs on certain interfaces in an STP, RSTP, or MSTP topology, however, can lead to network outages by triggering an STP misconfiguration. To prevent such outages, enable BPDU protection on those interfaces that should not receive BPDUs.

Enable BPDU protection on switch interfaces connected to user devices or on interfaces on which no BPDUs are expected, such as edge ports. If a BPDU is received on a BPDU-protected interface, the interface is disabled and stops forwarding frames.

Two J-EX Series switches are displayed in Figure 35 on page 1318. In this example, Switch 1 and Switch 2 are configured for RSTP and create a loop-free topology. The interfaces on Switch 2 are access ports.

This example shows you how to configure interface **ge-0/0/5** and interface **ge-0/0/6** as edge ports and to configure BPDU protection. When BPDU protection is enabled, the interfaces will transition to a blocking state when BPDUs are received on them.

Figure 35: BPDU Protection Topology

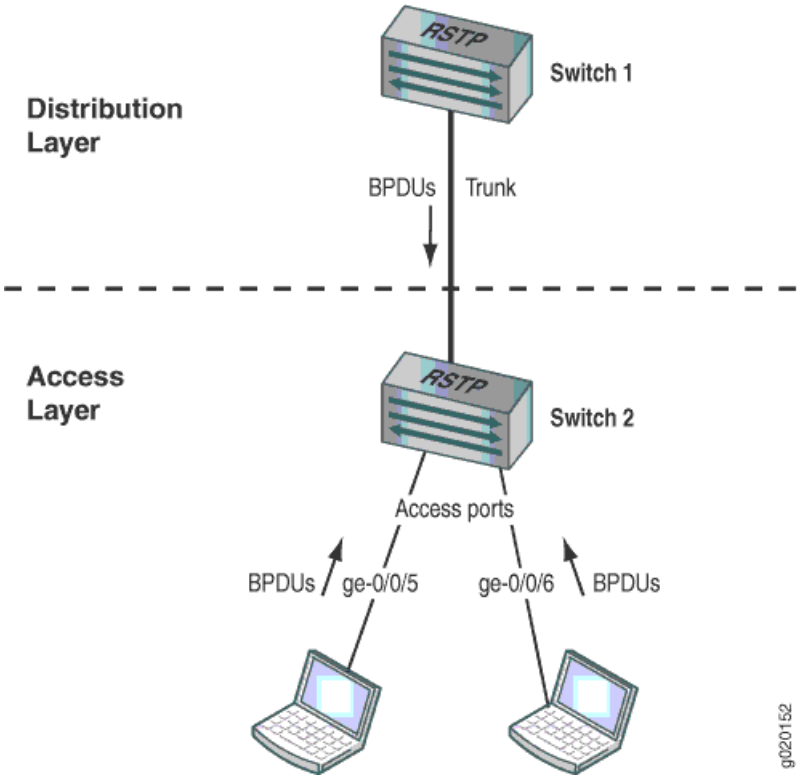

Table 173 on page 1318 shows the components that will be configured for BPDU protection.

Table 173: Components of the Topology for Configuring BPDU Protection on J-EX Series Switches

Property	Settings
Switch 1 (Distribution Layer)	Switch 1 is connected to Switch 2 on a trunk interface.
Switch 2 (Access Layer)	Switch 2 has these access ports that require BPDU protection: • **ge-0/0/5** • **ge-0/0/6**

This configuration example is using an RSTP topology. You also can configure BPDU protection for STP or MSTP topologies at the [**edit protocols** (**mstp** | **stp**)] hierarchy level.

Configuration

To configure BPDU protection on two access interfaces:

CLI Quick Configuration

To quickly configure BPDU protection on Switch 2, copy the following commands and paste them into the switch terminal window:

```
[edit]
set protocols rstp interface ge-0/0/5 edge
set protocols rstp interface ge-0/0/6 edge
set protocols rstp bpdu-block-on-edge
```

Step-by-Step Procedure

To configure BPDU protection:

1. Configure interface **ge-0/0/5** and interface **ge-0/0/6** on Switch 2 as edge ports:

```
[edit protocols rstp]
user@switch# set interface ge-0/0/5 edge
user@switch#set interface ge-0/0/6 edge
```

2. Configure BPDU protection on all edge ports:

```
[edit protocols rstp]
user@switch# set bpdu-block-on-edge
```

Results

Check the results of the configuration:

```
user@switch> show configuration protocols rstp
interface ge-0/0/5.0 {
  edge;
}
interface ge-0/0/6.0 {
  edge;
}
bpdu-block-on-edge;
```

Verification

To confirm that the configuration is working properly:

- Displaying the Interface State Before BPDU Protection Is Triggered on page 1319
- Verifying That BPDU Protection is Working Correctly on page 1320

Displaying the Interface State Before BPDU Protection Is Triggered

Purpose

Before BPDUs are being received from the PCs connected to interface **ge-0/0/5** and interface **ge-0/0/6**, confirm the interface state.

Action

Use the operational mode command:

```
user@switch>  show spanning-tree interface

Spanning tree interface parameters for instance 0

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ge-0/0/0.0     128:513      128:513  32768.0019e2503f00     20000  BLK    DIS
ge-0/0/1.0     128:514      128:514  32768.0019e2503f00     20000  BLK    DIS
ge-0/0/2.0     128:515      128:515  32768.0019e2503f00     20000  BLK    DIS
ge-0/0/3.0     128:516      128:516  32768.0019e2503f00     20000  FWD    DESG
ge-0/0/4.0     128:517      128:517  32768.0019e2503f00     20000  FWD    DESG
ge-0/0/5.0     128:518      128:518  32768.0019e2503f00     20000  FWD    DESG
ge-0/0/6.0     128:519      128:519  32768.0019e2503f00     20000  FWD    DESG
[output truncated]
```

Meaning

The output from the operational mode command **show spanning-tree interface** shows that **ge-0/0/5.0** and interface **ge-0/0/6.0** are designated ports in a forwarding state.

Verifying That BPDU Protection is Working Correctly

Purpose In this example, the PCs connected to Switch 2 start sending BPDUs to interface **ge-0/0/5.0** and interface **ge-0/0/6.0** . Verify that BPDU protection is configured on the interfaces.

Action Use the operational mode command:

```
user@switch> show spanning-tree interface

Spanning tree interface parameters for instance 0

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ge-0/0/0.0     128:513      128:513  32768.0019e2503f00     20000  BLK    DIS
ge-0/0/1.0     128:514      128:514  32768.0019e2503f00     20000  BLK    DIS
ge-0/0/2.0     128:515      128:515  32768.0019e2503f00     20000  BLK    DIS
ge-0/0/3.0     128:516      128:516  32768.0019e2503f00     20000  FWD    DESG
ge-0/0/4.0     128:517      128:517  32768.0019e2503f00     20000  FWD    DESG
ge-0/0/5.0     128:518      128:518  32768.0019e2503f00     20000  BLK    DIS
(Bpdu-Incon)
ge-0/0/6.0     128:519      128:519  32768.0019e2503f00     20000  BLK    DIS
(Bpdu-Incon)
ge-0/0/7.0     128:520        128:1  16384.00aabbcc0348     20000  FWD    ROOT
ge-0/0/8.0     128:521      128:521  32768.0019e2503f00     20000  FWD    DESG
[output truncated]
```

Meaning When BPDUs are sent from the PCs to interface **ge-0/0/5.0** and interface **ge-0/0/6.0** on Switch 2, the output from the operational mode command **show spanning-tree interface** shows that the interfaces have transitioned to a BPDU inconsistent state. The BPDU inconsistent state makes the interfaces block and prevents them from forwarding traffic.

Disabling the BPDU protection configuration on an interface does not unblock the interface. If the **disable-timeout** statement has been included in the BPDU configuration, the interface automatically returns to service after the timer expires. Otherwise, use the operational mode command **clear ethernet-switching bpdu-error** to unblock the interface.

If the PCs connected to Switch 2 send BPDUs to the interfaces again, BPDU protection is triggered once more and the interfaces transition back to the BPDU inconsistent state. In such cases, you need to find and repair the misconfiguration on the PCs that is triggering BPDUs being sent to Switch 2.

Related Documentation

- Example: Configuring Faster Convergence and Improving Network Stability with RSTP on J-EX Series Switches on page 1283
- Example: Configuring BPDU Protection on non-STP Interfaces to Prevent STP Miscalculations on J-EX Series Switches on page 1321
- Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on J-EX Series Switches on page 1325
- Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees on J-EX Series Switches on page 1329

- Understanding BPDU Protection for STP, RSTP, and MSTP on J-EX Series Switches on page 1278

Example: Configuring BPDU Protection on non-STP Interfaces to Prevent STP Miscalculations on J-EX Series Switches

J-EX Series switches provide Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP). Configure BPDU protection on non-STP interfaces that are connected to switches with spanning trees to prevent the non-STP interfaces from receiving BPDUs. When non-STP interfaces receive BPDUs, it can result in an STP misconfiguration, which could lead to network outages.

This example describes how to configure BPDU protection on non-STP interfaces on a J-EX Series switch:

- Requirements on page 1321
- Overview and Topology on page 1321
- Configuration on page 1323
- Verification on page 1323

Requirements

This example uses the following hardware and software components:

- One J-EX Series switch in an RSTP topology
- One J-EX Series switch that is not in a spanning-tree topology

Before you configure the interface for BPDU protection, be sure you have:

- RSTP operating on Switch 1.
- Disabled RSTP on Switch 2.

NOTE: By default, RSTP is enabled on all J-EX Series switches.

Overview and Topology

A loop-free network is supported through the exchange of a special type of frame called bridge protocol data unit (BPDU). Receipt of BPDUs on certain interfaces can lead to network outages by triggering an STP miscalculation. Enable BPDU protection on those interfaces that should not receive BPDUs to prevent network outages.

BPDU protection for non-STP interfaces can be enabled on interfaces on a non-STP switch connected to an STP switch through a trunk interface. Enable BPDU protection on interfaces on which no BPDUs are expected, such as access ports connected to user devices. If BPDUs are received on a BPDU-protected interface, the interface transitions to a blocking state and stops forwarding frames.

Two J-EX Series switches are displayed in Figure 36 on page 1322. In this example, Switch 1 and Switch 2 are connected through a trunk interface. Switch 1 is configured for RSTP, but Switch 2 has no spanning tree. Switch 2 has two access ports: interface **ge-0/0/5** and interface **ge-0/0/6**.

This example shows you how to configure BPDU protection on interface **ge-0/0/5** and interface **ge-0/0/6**. When BPDU protection is enabled, the interfaces will transition to a blocking state if BPDUs are received.

Figure 36: BPDU Protection Topology

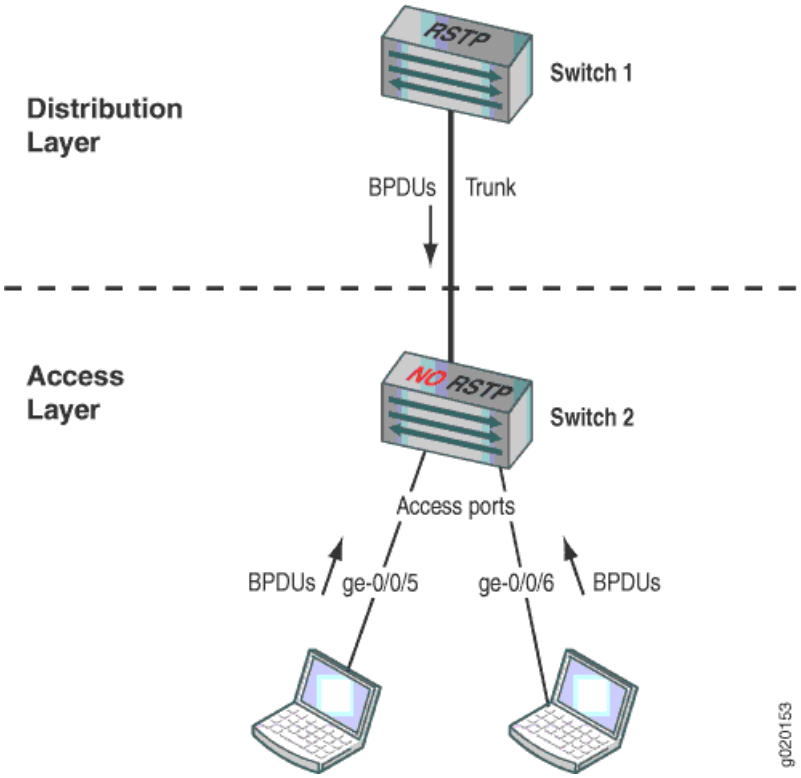

Table 174 on page 1322 shows the components that will be configured for BPDU protection.

Table 174: Components of the Topology for Configuring BPDU Protection on J-EX Series Switches

Property	Settings
Switch 1 (Distribution Layer)	Switch 1 is connected to Switch 2 through a trunk interface. Switch 1 is configured for RSTP.
Switch 2 (Access Layer)	Switch 2 has RSTP disabled and has these access ports that require BPDU protection: • **ge-0/0/5** • **ge-0/0/6**

CAUTION: When configuring BPDU protection on a non-STP configured switch connected to an STP-configured switch, be careful that you do not configure BPDU protection on all interfaces. Doing so could prevent BPDUs being received on interfaces (such as a trunk interface) that should be receiving BPDUs from an STP-configured switch.

Configuration

To configure BPDU protection on the interfaces:

CLI Quick Configuration

To quickly configure BPDU protection on Switch 2, copy the following commands and paste them into the switch terminal window:

```
[edit]
set ethernet-switching-options bpdu-block interface ge-0/0/5
set ethernet-switching-options bpdu-block interface ge-0/0/6
```

Step-by-Step Procedure

To configure BPDU protection:

1. Configure interface **ge-0/0/5** and interface **ge-0/0/6** on Switch 2:

```
[edit ethernet-switching-options]
user@switch# set bpdu-block interface ge-0/0/5
user@switch# set bpdu-block interface ge-0/0/6
```

Results

Check the results of the configuration:

```
user@switch> show ethernet-switching-options
bpdu-block {
   interface ge-0/0/5.0;
   interface ge-0/0/6.0;
}
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Displaying the Interface State Before BPDU Protection Is Triggered on page 1323
- Verifying That BPDU Protection Is Working Correctly on page 1324

Displaying the Interface State Before BPDU Protection Is Triggered

Purpose

Before BPDUs are being received from the PCs connected to interface **ge-0/0/5** and interface **ge-0/0/6**, confirm the interface state.

Action

Use the operational mode command:

```
user@switch> show ethernet-switching interfaces

Interface   State    VLAN members          Blocking
ge-0/0/0.0  down     default               unblocked
ge-0/0/1.0  down     default               unblocked
ge-0/0/2.0  down     default               unblocked
ge-0/0/3.0  up       default               unblocked
```

```
ge-0/0/4.0  up       v1                   unblocked
ge-0/0/5.0  up       v1                   unblocked
ge-0/0/6.0  up       default              unblocked
[output truncated]
```

Meaning The output from the operational mode command **show ethernet-switching interfaces** shows that **ge-0/0/5.0** and interface **ge-0/0/6.0** are **up** and unblocked.

Verifying That BPDU Protection Is Working Correctly

Purpose In this example, the PCs connected to Switch 2 start sending BPDUs to interface **ge-0/0/5.0** and interface **ge-0/0/6.0**. Verify that BPDU protection is configured on the interfaces.

Action Use the operational mode command:

```
user@switch> show ethernet-switching interfaces

Interface   State    VLAN members          Blocking
ge-0/0/0.0  up       default               unblocked
ge-0/0/1.0  up       default               unblocked
ge-0/0/2.0  up       default               unblocked
ge-0/0/3.0  up       default               unblocked
ge-0/0/4.0  up       v1                    unblocked
ge-0/0/5.0  down     v1                    blocked - blocked by bpdu-control
ge-0/0/6.0  down     default               blocked - blocked by bpdu-control
[output truncated]
```

Meaning When BPDUs are sent from the PCs to interface **ge-0/0/5.0** and interface **ge-0/0/6.0** on Switch 2, the output from the operational mode command **show spanning-tree interface** shows that the interfaces have transitioned to a BPDU inconsistent state. The BPDU inconsistent state makes the interfaces shut down and prevents them from forwarding traffic.

Disabling the BPDU protection configuration on an interface does not unblock the interface. If the **disable-timeout** statement has been included in the BPDU configuration, the interface automatically returns to service after the timer expires. Otherwise, use the operational mode command **clear ethernet-switching bpdu-error** to recover from the error condition and restore the interface to service.

If the PCs connected to Switch 2 send BPDUs to the interfaces again, BPDU protection is triggered once more and the interfaces transition back to the BPDU inconsistent state. In such cases, you need to find and repair the misconfiguration on the PCs that is triggering BPDUs being sent to Switch 2.

Related Documentation

- Example: Configuring Faster Convergence and Improving Network Stability with RSTP on J-EX Series Switches on page 1283
- Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations on J-EX Series Switches on page 1317
- Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on J-EX Series Switches on page 1325

- Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees on J-EX Series Switches on page 1329
- Understanding BPDU Protection for STP, RSTP, and MSTP on J-EX Series Switches on page 1278

Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on J-EX Series Switches

J-EX Series switches provide Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP). Loop protection increases the efficiency of STP, RSTP, and MSTP by preventing interfaces from moving into a forwarding state that would result in a loop opening up in the network.

This example describes how to configure loop protection for an interface on a J-EX Series switch in an RSTP topology:

- Requirements on page 1325
- Overview and Topology on page 1325
- Configuration on page 1327
- Verification on page 1327

Requirements

This example uses the following hardware and software components:

- Three J-EX Series switches in an RSTP topology

Before you configure the interface for loop protection, be sure you have:

- RSTP operating on the switches.

NOTE: By default, RSTP is enabled on all J-EX Series switches.

Overview and Topology

A loop-free network in spanning-tree topologies is supported through the exchange of a special type of frame called bridge protocol data unit (BPDU). Peer STP applications running on the switch interfaces use BPDUs to communicate. Ultimately, the exchange of BPDUs determines which interfaces block traffic (preventing loops) and which interfaces become root ports and forward traffic.

A blocking interface can transition to the forwarding state in error if the interface stops receiving BPDUs from its designated port on the segment. Such a transition error can occur when there is a hardware error on the switch or software configuration error between the switch and its neighbor. When this happens, a loop opens up in the spanning tree. Loops in a Layer 2 topology cause broadcast, unicast, and multicast frames to

continuously circle the looped network. As a switch processes a flood of frames in a looped network, its resources become depleted and the ultimate result is a network outage.

CAUTION: An interface can be configured for either loop protection or root protection, but not for both.

Three J-EX Series switches are displayed in Figure 37 on page 1326. In this example, they are configured for RSTP and create a loop-free topology. Interface **ge-0/0/6** is blocking traffic between Switch 3 and Switch 1; thus, traffic is forwarded through interface **ge-0/0/7** on Switch 2. BPDUs are being sent from the root bridge on Switch 1 to both of these interfaces.

This example shows how to configure loop protection on interface **ge-0/0/6** to prevent it from transitioning from a blocking state to a forwarding state and creating a loop in the spanning-tree topology.

Figure 37: Network Topology for Loop Protection

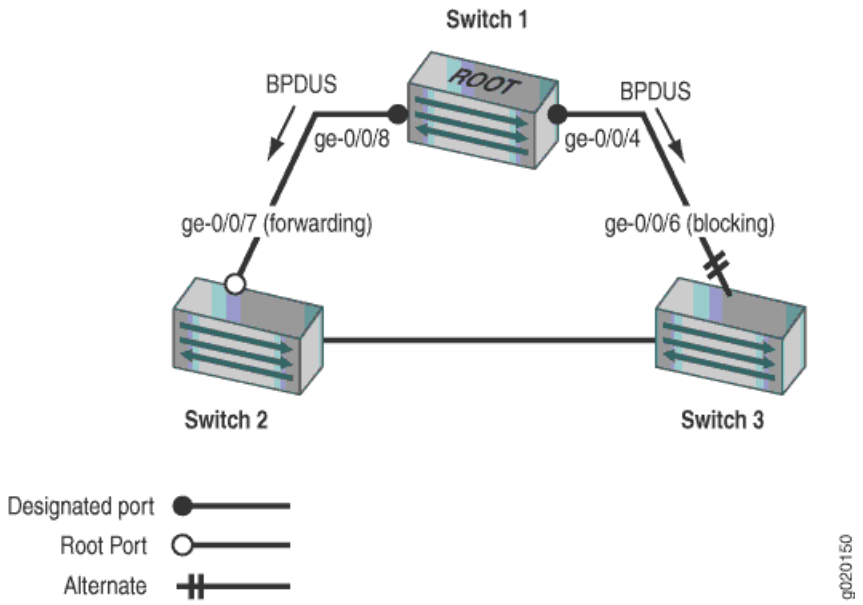

Table 175 on page 1326 shows the components that will be configured for loop protection.

Table 175: Components of the Topology for Configuring Loop Protection on J-EX Series Switches

Property	Settings
Switch 1	Switch 1 is the root bridge.
Switch 2	Switch 2 has the root port **ge-0/0/7**.
Switch 3	Switch 3 is connected to Switch 1 through interface **ge-0/0/6**.

A spanning-tree topology contains ports that have specific roles:

- The root port is responsible for forwarding data to the root bridge.
- The alternate port is a standby port for the root port. When a root port goes down, the alternate port becomes the active root port.
- The designated port forwards data to the downstream network segment or device.

This configuration example uses an RSTP topology. However, you also can configure loop protection for STP or MSTP topologies at the [**edit protocols** (**mstp** | **stp**)] hierarchy level.

Configuration

To configure loop protection on an interface:

CLI Quick Configuration

To quickly configure loop protection on interface **ge-0/0/6**:

```
[edit]
set protocols rstp interface ge-0/0/6 bpdu-timeout-action block
```

Step-by-Step Procedure

To configure loop protection:

1. Configure interface **ge-0/0/6** on Switch 3:

```
[edit protocols rstp]
user@switch# set interface ge-0/0/6 bpdu-timeout-action block
```

Results

Check the results of the configuration:

```
user@switch> show configuration protocols rstp
interface ge-0/0/6.0 {
  bpdu-timeout-action {
    block;
  }
}
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Displaying the Interface State Before Loop Protection Is Triggered on page 1327
- Verifying That Loop Protection Is Working on an Interface on page 1328

Displaying the Interface State Before Loop Protection Is Triggered

Purpose

Before loop protection is triggered on interface **ge-0/0/6**, confirm that the interface is blocking.

Action

Use the operational mode command:

```
user@switch> show spanning-tree interface

Spanning tree interface parameters for instance 0

Interface    Port ID    Designated      Designated         Port    State  Role
```

```
                        port ID        bridge ID          Cost
ge-0/0/0.0     128:513     128:513  32768.0019e2503f00     20000  BLK    DIS
ge-0/0/1.0     128:514     128:514  32768.0019e2503f00     20000  BLK    DIS
ge-0/0/2.0     128:515     128:515  32768.0019e2503f00     20000  BLK    DIS
ge-0/0/3.0     128:516     128:516  32768.0019e2503f00     20000  FWD    DESG
ge-0/0/4.0     128:517     128:517  32768.0019e2503f00     20000  FWD    DESG
ge-0/0/5.0     128:518     128:518  32768.0019e2503f00     20000  FWD    DESG
ge-0/0/6.0     128:519       128:2  16384.00aabbcc0348     20000  BLK    ALT
[output truncated]
```

Meaning The output from the operational mode command **show spanning-tree interface** shows that **ge-0/0/6.0** is the alternate port and in a blocking state.

Verifying That Loop Protection Is Working on an Interface

Purpose Verify the loop protection configuration on interface **ge-0/0/6**. RSTP has been disabled on interface **ge-0/0/4** on Switch 1. This will stop BPDUs from being sent to interface **ge-0/0/6** and trigger loop protection on the interface.

Action Use the operational mode command:

```
user@switch> show spanning-tree interface

Spanning tree interface parameters for instance 0

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ge-0/0/0.0     128:513     128:513  32768.0019e2503f00     20000  BLK    DIS
ge-0/0/1.0     128:514     128:514  32768.0019e2503f00     20000  BLK    DIS
ge-0/0/2.0     128:515     128:515  32768.0019e2503f00     20000  BLK    DIS
ge-0/0/3.0     128:516     128:516  32768.0019e2503f00     20000  FWD    DESG
ge-0/0/4.0     128:517     128:517  32768.0019e2503f00     20000  FWD    DESG
ge-0/0/5.0     128:518     128:518  32768.0019e2503f00     20000  FWD    DESG
ge-0/0/6.0     128:519     128:519  32768.0019e2503f00     20000  BLK    DIS
(Loop-Incon)
[output truncated]
```

Meaning The operational mode command **show spanning-tree interface** shows that interface **ge-0/0/6.0** has detected that BPDUs are no longer being forwarded to it and has moved into a loop-inconsistent state. The loop-inconsistent state prevents the interface from transitioning to a forwarding state. The interface recovers and transitions back to its original state as soon as it receives BPDUs.

Related Documentation

- Example: Configuring Faster Convergence and Improving Network Stability with RSTP on J-EX Series Switches on page 1283
- Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees on J-EX Series Switches on page 1329
- Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations on J-EX Series Switches on page 1317
- Example: Configuring BPDU Protection on non-STP Interfaces to Prevent STP Miscalculations on J-EX Series Switches on page 1321

Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees on J-EX Series Switches

J-EX Series switches provide Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree protocol (RSTP), and Multiple Spanning Tree Protocol (MSTP). Root protection increases the efficiency of STP, RSTP, and MSTP by allowing network administrators to manually enforce the root bridge placement in the network.

This example describes how to configure root protection on an interface on a J-EX Series switch:

Requirements

This example uses the following hardware and software components:

- Four J-EX Series switches in an RSTP topology

Before you configure the interface for root protection, be sure you have:

- RSTP operating on the switches.

NOTE: By default, RSTP is enabled on all J-EX Series switches.

Overview and Topology

Peer STP applications running on switch interfaces exchange a special type of frame called a bridge protocol data unit (BPDU). Switches communicate interface information using BPDUs to create a loop-free topology that ultimately determines the root bridge and which interfaces block or forward traffic in the spanning tree.

However, a root port elected through this process has the possibility of being wrongly elected. A user bridge application running on a PC can generate BPDUs, too, and interfere with root port election.

To prevent this from happening, enable root protection on interfaces that should not receive superior BPDUs from the root bridge and should not be elected as the root port. These interfaces are typically located on an administrative boundary and are designated ports.

When root protection is enabled on an interface:

- The interface is blocked from becoming the root port.
- Root protection is enabled for all STP instances on that interface.
- The interface is blocked only for instances for which it receives superior BPDUs. Otherwise, it participates in the spanning-tree topology.

CAUTION: An interface can be configured for either root protection or loop protection, but not for both.

Four J-EX Series switches are displayed in Figure 38 on page 1330. In this example, they are configured for RSTP and create a loop-free topology. Interface **ge-0/0/7** on Switch 1 is a designated port on an administrative boundary. It connects to Switch 4. Switch 3 is the root bridge. Interface **ge-0/0/6** on Switch 1 is the root port.

This example shows how to configure root protection on interface **ge-0/0/7** to prevent it from transitioning to become the root port.

Figure 38: Network Topology for Root Protection

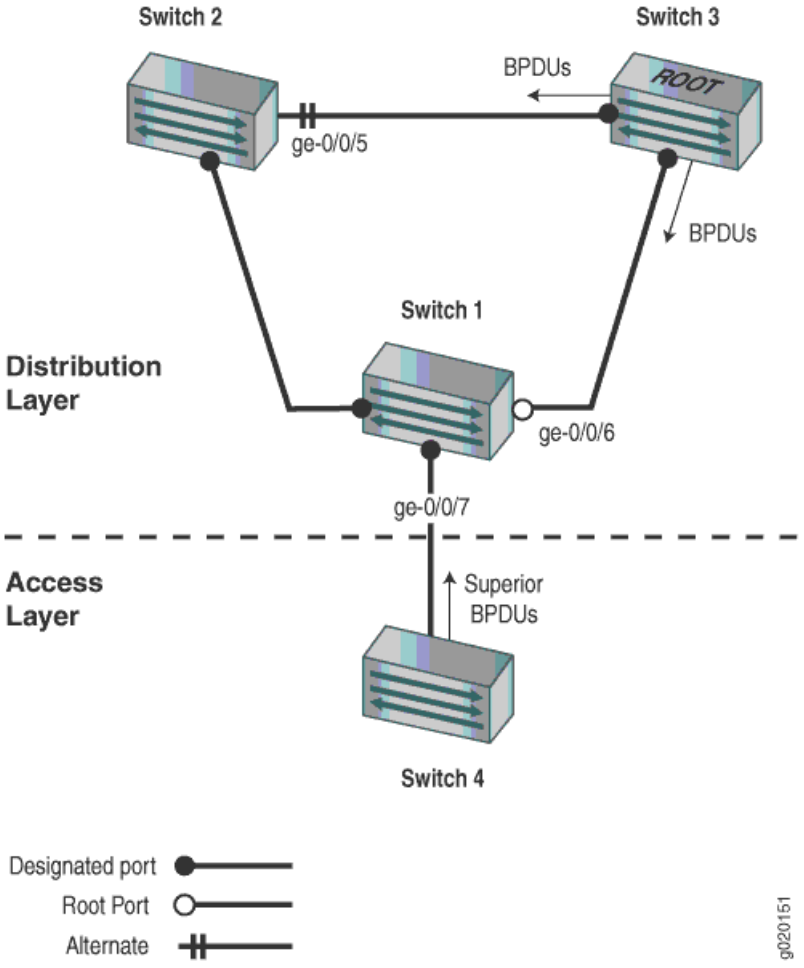

Table 176 on page 1331 shows the components that will be configured for root protection.

Table 176: Components of the Topology for Configuring Root Protection on J-EX Series Switches

Property	Settings
Switch 1	Switch 1 is connected to Switch 4 through interface **ge-0/0/7**.
Switch 2	Switch 2 is connected to Switch 1 and Switch 3. Interface **ge-0/0/4** is the alternate port in the RSTP topology.
Switch 3	Switch 3 is the root bridge and is connected to Switch 1 and Switch 2.
Switch 4	Switch 4 is connected to Switch 1. After loop protection is configured on interface **ge-0/0/7**, Switch 4 will send superior BPDUs that will trigger loop protection on interface **ge-0/0/7**.

A spanning tree topology contains ports that have specific roles:

- The root port is responsible for forwarding data to the root bridge.
- The alternate port is a standby port for the root port. When a root port goes down, the alternate port becomes the active root port.
- The designated port forwards data to the downstream network segment or device.

This configuration example uses an RSTP topology. However, you also can configure root protection for STP or MSTP topologies at the [**edit protocols** (**mstp** | **stp**)] hierarchy level.

Configuration

To configure root protection on an interface:

CLI Quick Configuration

To quickly configure root protection on interface **ge-0/0/7**, copy the following command and paste it into the switch terminal window:

```
[edit]
set protocols rstp interface ge-0/0/7 no-root-port
```

Step-by-Step Procedure

To configure root protection:

1. Configure interface **ge-0/0/7**:

```
[edit protocols rstp]

user@switch#
set interface ge-0/0/7 no-root-port
```

Results

Check the results of the configuration:

```
user@switch> show configuration protocols rstp
interface ge-0/0/7.0 {
   no-root-port;
}
```

Verification

To confirm that the configuration is working properly:

Displaying the Interface State Before Root Protection Is Triggered

Purpose Before root protection is triggered on interface **ge-0/0/7**, confirm the interface state.

Action Use the operational mode command:

```
user@switch>  show spanning-tree interface

Spanning tree interface parameters for instance 0

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ge-0/0/0.0     128:513      128:513  32768.0019e2503f00     20000  BLK    DIS
ge-0/0/1.0     128:514      128:514  32768.0019e2503f00     20000  BLK    DIS
ge-0/0/2.0     128:515      128:515  32768.0019e2503f00     20000  BLK    DIS
ge-0/0/3.0     128:516      128:516  32768.0019e2503f00     20000  FWD    DESG
ge-0/0/4.0     128:517      128:517  32768.0019e2503f00     20000  FWD    DESG
ge-0/0/5.0     128:518        128:2  16384.00aabbcc0348     20000  BLK    ALT
ge-0/0/6.0     128:519        128:1  16384.00aabbcc0348     20000  FWD    ROOT
ge-0/0/7.0     128:520      128:520  32768.0019e2503f00     20000  FWD    DESG
[output truncated]
```

Meaning The output from the operational mode command **show spanning-tree interface** shows that **ge-0/0/7.0** is a designated port in a forwarding state.

Verifying That Root Protection Is Working on the Interface

Purpose A configuration change takes place on Switch 4. A smaller bridge priority on the Switch 4 causes it to send superior BPDUs to interface **ge-0/0/7**. Receipt of superior BPDUs on interface **ge-0/0/7** will trigger root protection. Verify that root protection is operating on interface **ge-0/0/7**.

Action Use the operational mode command:

```
user@switch>  show spanning-tree interface

Spanning tree interface parameters for instance 0

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ge-0/0/0.0     128:513      128:513  32768.0019e2503f00     20000  BLK    DIS
ge-0/0/1.0     128:514      128:514  32768.0019e2503f00     20000  BLK    DIS
ge-0/0/2.0     128:515      128:515  32768.0019e2503f00     20000  BLK    DIS
ge-0/0/3.0     128:516      128:516  32768.0019e2503f00     20000  FWD    DESG
ge-0/0/4.0     128:517      128:517  32768.0019e2503f00     20000  FWD    DESG
ge-0/0/5.0     128:518        128:2  16384.00aabbcc0348     20000  BLK    ALT
ge-0/0/6.0     128:519        128:1  16384.00aabbcc0348     20000  FWD    ROOT
ge-0/0/7.0     128:520      128:520  32768.0019e2503f00     20000  BLK    DIS
```

```
(Root–Incon)
[output truncated]
```

Meaning The operational mode command **show spanning-tree interface** shows that interface **ge-0/0/7.0** has transitioned to a loop inconsistent state. The loop inconsistent state makes the interface block and prevents the interface from becoming a candidate for the root port. When the root bridge no longer receives superior STP BPDUs from the interface, the interface will recover and transition back to a forwarding state. Recovery is automatic.

Related Documentation

- Example: Configuring Faster Convergence and Improving Network Stability with RSTP on J-EX Series Switches on page 1283
- Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on J-EX Series Switches on page 1325
- Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations on J-EX Series Switches on page 1317
- Example: Configuring BPDU Protection on non-STP Interfaces to Prevent STP Miscalculations on J-EX Series Switches on page 1321
- Understanding Root Protection for STP, RSTP, VSTP, and MSTP on J-EX Series Switches on page 1280

CHAPTER 66

Configuring Spanning-Tree Protocols

Unblocking an Interface That Receives BPDUs in Error (CLI Procedure)

J-EX Series switches use bridge protocol data unit (BPDU) protection on interfaces to prevent them from receiving BPDUs that could trigger a spanning-tree misconfiguration. If BPDUs are received on a BPDU-protected interface, the interface transitions to a blocking state and stops forwarding frames.

After the misconfiguration that triggered the BPDUs being sent to an interface is fixed in the topology, the interface can be unblocked and returned to service.

To unblock an interface and return it to service using the CLI:

- Automatically unblock an interface by configuring a timer that expires (here, the interface is **ge-0/0/6**):

```
[edit ethernet-switching-options]
user@switch# set bpdu-block disable-timeout 30 interface ge-0/0/6
```

- Manually unblock an interface using the operational mode command:

```
user@switch> clear ethernet-switching bpdu-error interface ge-0/0/6
```

Related Documentation

- Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations on J-EX Series Switches on page 1317
- Example: Configuring BPDU Protection on non-STP Interfaces to Prevent STP Miscalculations on J-EX Series Switches on page 1321
- Understanding BPDU Protection for STP, RSTP, and MSTP on J-EX Series Switches on page 1278

Configuring STP (CLI Procedure)

The default spanning-tree protocol for J-EX Series switches is Rapid Spanning Tree Protocol (RSTP). RSTP provides faster convergence times than Spanning Tree Protocol (STP). However, some legacy networks require the slower convergence times of basic STP.

If your network includes 802.1D 1998 bridges, you can remove RSTP and explicitly configure STP. When you explicitly configure STP, the J-EX Series switches use the IEEE 802.1D 2004 specification, force version 0. This configuration runs a version of RSTP that is compatible with the classic, basic STP.

To configure STP using the CLI:

1. Delete the RSTP configuration on the interface (here, the interface is **ge-0/0/5**):

```
[edit]
user@switch# delete protocols rstp interface ge-0/0/5
```

2. Configure STP on the interface:

```
[edit]
user@switch# set protocols stp interface ge-0/0/5
```

3. Commit the configuration:

```
[edit]
user@switch# commit
```

Related Documentation

- **show spanning-tree bridge on page 1398**
- **show spanning-tree interface on page 1407**
- Understanding STP for J-EX Series Switches on page 1275

Configuring Spanning-Tree Protocols (J-Web Procedure)

J-EX Series switches provide Layer 2 loop prevention through Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), Multiple Spanning Tree Protocol (MSTP), and VLAN Spanning Tree Protocol (VSTP). You can configure STP, RSTP, and MSTP using the J-Web interface. You can configure bridge protocol data unit (BPDU) protection on interfaces to prevent them from receiving BPDUs that could result in STP misconfigurations, which could lead to network outages.

To configure STP, MSTP, or RSTP for a J-EX Series switch using the J-Web interface:

1. Select **Configure** > **Switching** > **Spanning Tree**.

   The Spanning Tree Configuration page displays the spanning-tree protocol configuration parameters and a list of interfaces configured for each spanning-tree protocol configuration.

NOTE: After you make changes to the configuration in this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See "Using the Commit Options to Commit Configuration Changes (J-Web Procedure)" on page 334 for details about all commit options.

2. Click one:
   - **Add**—Creates a spanning-tree protocol configuration.
     a. Select a protocol name.
     b. Enter information as described in Table 177 on page 1337.
     c. Click **OK** to apply changes to the configuration or click **Cancel** to cancel without saving changes.
   - **Edit**—Modifies a selected spanning-tree protocol configuration.
     a. Enter information as described in Table 177 on page 1337.
     b. Click **OK** to apply changes to the configuration or click **Cancel** to cancel without saving changes.
   - **Delete**—Deletes a selected spanning-tree protocol configuration.

Table 177: Spanning-Tree Protocol Configuration Parameters

Field	Function	Your Action
General		
Protocol Name	Specifies the spanning-tree protocol type: STP, MSTP, or RSTP.	None.
Disable	Disables spanning-tree protocol on the interface.	To enable this option, select the check box.
BPDU Protect	Specifies BPDU protection on all edge interfaces on the switch.	To enable this option, select the check box.
Bridge Priority	Specifies the bridge priority. The bridge priority determines which bridge is elected as the root bridge. If two bridges have the same path cost to the root bridge, the bridge priority determines which bridge becomes the designated bridge for a LAN segment.	Select a value from the list.
Forward Delay	Specifies the number of seconds an interface waits before changing from spanning-tree learning and listening states to the forwarding state.	Type a value.
Hello Time	Specifies the time interval in seconds at which the root bridge transmits configuration BPDUs.	Type a value.

Table 177: Spanning-Tree Protocol Configuration Parameters *(continued)*

Field	Function	Your Action
Max Age	Specifies the maximum-aging time in seconds for all MST instances. The maximum aging time is the number of seconds a switch waits without receiving spanning-tree configuration messages before attempting a reconfiguration.	Type a value.
Max Hops	(MSTP only) Specifies the number of hops in a region before the BPDU is discarded.	Type a value.
Configuration Name	(MSTP only) Specifies the MSTP region name carried in the MSTP BPDUs.	Type a name.
Revision Level	(MSTP only) Specifies the revision number of the MSTP configuration.	Type a value.
Ports		
Interface Name	Specifies an interface for the spanning-tree protocol.	1. Click the **Ports** tab. 2. Choose one: • Click **Add** and select an interface from the list. • Select an interface in the **Port/State** table and click **Edit**. • To delete an interface from the configuration, select it in the **Port/State** table and click **Remove**.
Cost	Specifies the link cost to determine which bridge is the designated bridge and which interface is the designated interface.	Type a value.
Priority	Specifies the interface priority to determine which interface is elected as the root port.	Select a value from the list.
Disable Port	Disables the spanning-tree protocol on the interface.	To enable the option, select the check box.
Edge	Configures the interface as an edge interface. Edge interfaces immediately transition to a forwarding state.	To enable the option, select the check box.
No Root Port	Specifies an interface as a spanning-tree designated port. If the bridge receives superior STP BPDUs on a root-protected interface, that interface transitions to a root-prevented STP state (inconsistency state) and the interface is blocked. This blocking prevents a bridge that should not be the root bridge from being elected the root bridge. When the bridge stops receiving superior STP BPDUs on the root-protected interface, interface traffic is no longer blocked.	To enable the option, select the check box.

Table 177: Spanning-Tree Protocol Configuration Parameters *(continued)*

Field	Function	Your Action
Interface Mode	Specifies the link mode.	1. To enable the option, select the check box. 2. Select one: • **Point to Point**—For a full-duplex link, the default link mode is point-to-point. • **Shared**—For a half-duplex link, the default link mode is shared.
BPDU Timeout Action	Specifies the BPDU timeout action for the interface.	Select one: • **Alarm** • **Block**
MSTI (MSTP only)		
MSTI Name	Specifies a name (an MSTI ID) for the MST instance.	1. Click the **MSTI** tab. 2. Choose one: • Click **Add**. • Select an MSTI ID and click **Edit**. • To delete an MSTI from the configuration, select the MSTI ID and slick **Remove**.
Bridge Priority	Specifies the bridge priority. The bridge priority determines which bridge is elected as the root bridge. If two bridges have the same path cost to the root bridge, the bridge priority determines which bridge becomes the designated bridge for a LAN segment.	Select a value from the list.
VLAN ID	Specifies the VLAN for the MST instance.	In the VLAN box, choose one: • Click **Add**, select a VLAN from the list and click **OK**. • To remove a VLAN association, select the VLAN ID, click **Remove**, and click **OK**.

Table 177: Spanning-Tree Protocol Configuration Parameters *(continued)*

Field	Function	Your Action
Interfaces	Specifies an interface for the MST instance.	1. In the Interfaces box, click **Add** and select an interface from the list, or select an interface from the list and click **Edit**. 2. Specify the link cost to determine which bridge is the designated bridge and which interface is the designated interface. 3. Specify the interface priority to determine which interface is elected as the root port. 4. If you want to disable the interface, select the check box. 5. Click OK. To delete an interface configuration, select the interface, click **Remove**, and click **OK**.

Related Documentation

- Configuring STP (CLI Procedure) on page 1336
- Monitoring Spanning-Tree Protocols on page 1343
- Unblocking an Interface That Receives BPDUs in Error (CLI Procedure) on page 1335
- Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations on J-EX Series Switches on page 1317
- Example: Configuring Network Regions for VLANs with MSTP on J-EX Series Switches on page 1297
- Example: Configuring Faster Convergence and Improving Network Stability with RSTP on J-EX Series Switches on page 1283

Configuring VLAN Spanning Tree Protocol (CLI Procedure)

VLAN Spanning Tree Protocol (VSTP) allows J-EX Series switches to run one or more Spanning Tree Protocol (STP) or Rapid Spanning Tree Protocol (RSTP) instances for each VLAN on which VSTP is enabled. For networks with multiple VLANs, VSTP improves intelligent tree spanning by defining best paths within the VLANs instead of within the entire network.

To configure VSTP:

1. (Optional) Enable Rapid Spanning Tree Protocol (RSTP):

```
[edit protocols]
user@switch# set rstp
```

   VSTP can run on a maximum of 253 VLANs; RSTP runs on the remaining VLANs if configured. Enabling RSTP ensures that a spanning-tree protocol runs on all VLANs.

2. Enable VSTP.

- To enable VSTP on multiple VLANs using a VLAN group:

```
[edit protocols]
user@switch# set vstp vlan-group group group-name vlan vlan-id-range
```

- To enable VSTP on all VLANs:

```
[edit protocols]
user@switch# set vstp vlan all
```

NOTE: RSTP must be enabled if the set vstp vlan all statement is used to enable VSTP and the switch has more than 253 VLANs. If the set vstp vlan all statement is used to enable VSTP on a switch with more than 253 VLANs, the configuration cannot be committed.

- To enable VSTP on a VLAN using a single VLAN ID:

```
[edit protocols]
user@switch# set vstp vlan vlan-id
```

- To enable VSTP on a VLAN using a single VLAN name:

```
[edit protocols]
user@switch# set vstp vlan vlan-name
```

Related Documentation

- Understanding VSTP for J-EX Series Switches on page 1281

CHAPTER 67

Verifying Spanning Tree Protocols

Monitoring Spanning-Tree Protocols

Purpose Use the monitoring feature to view status and information about the spanning-tree protocol parameters on your J-EX Series switch.

Action To display spanning-tree protocol parameter details in the J-Web interface, select **Monitor > Switching > STP**.

To display spanning-tree protocol parameter details in the CLI, enter the following commands:

- **show spanning-tree interface**
- **show spanning-tree bridge**

Meaning Table 178 on page 1343 summarizes the spanning-tree protocol parameters.

Table 178: Summary of Spanning-Tree Protocols Output Fields

Field	Values
Bridge Parameters	
Context ID	An internally generated identifier.
Enabled Protocol	Spanning-tree protocol type enabled.
Root ID	Bridge ID of the elected spanning-tree root bridge. The bridge ID consists of a configurable bridge priority and the MAC address of the bridge.
Bridge ID	Locally configured bridge ID.
Hello Time	The time for which the bridge interface remains in the listening or learning state.
Forward Delay	The time for which the bridge interface remains in the listening or learning state before transitioning to the forwarding state.

Table 178: Summary of Spanning-Tree Protocols Output Fields *(continued)*

Field	Values
Extended System ID	The system ID.
Inter Instance ID	An internally generated instance identifier.
Maximum Age	Maximum age of received bridge protocol data units (BPDUs).
Number of topology changes	Total number of STP topology changes detected since the switch last booted.
Spanning Tree Interface Details	
Interface Name	Interface configured to participate in the STP instance.
Port ID	Logical interface identifier configured to participate in the STP instance.
Designated Port ID	Port ID of the designated port for the LAN segment to which the interface is attached.
Designated Bridge ID	ID of the designated bridge to which the interface is attached.
Port Cost	Configured cost for the interface.
Port State	STP port state: • Forwarding (FWD) • Blocking (BLK) • Listening • Learning • Disabled
Role	MSTP or RSTP port role, Designated (DESG), backup (BKUP), alternate (ALT), or root.
Spanning Tree Statistics of Interface	
Interface	Interface for which statistics is being displayed.
BPDUs Sent	Total number of BPDUs sent.
BPDUs Received	Total number of BPDUs received.
Next BPDU Transmission	Number of seconds until the next BPDU is scheduled to be sent.

Related Documentation

- **show spanning-tree interface on page 1407**
- **show spanning-tree bridge on page 1398**
- Configuring Spanning-Tree Protocols (J-Web Procedure) on page 1336
- Configuring STP (CLI Procedure) on page 1336

- Example: Configuring Network Regions for VLANs with MSTP on J-EX Series Switches on page 1297
- Example: Configuring Faster Convergence and Improving Network Stability with RSTP on J-EX Series Switches on page 1283

CHAPTER 68

Configuration Statements for Spanning-Tree Protocols

[edit protocols] Configuration Statement Hierarchy

```
protocols {
  connections {
    remote-interface-switch connection-name {
      interface interface-name.unit-number;
      transmit-lsp label-switched-path;
      receive-lsp label-switched-path;
    }
  }
  dot1x {
    authenticator {
      authentication-profile-name profile-name;
      interface (all | [ interface-names ]) {
        disable;
        guest-vlan ( vlan-id | vlan-name);
        mac-radius <restrict>;
        maximum-requests number;
        no-reauthentication;
        quiet-period seconds;
        reauthentication {
          interval seconds;
        }
        retries number;
        server-fail (deny | permit | use-cache | vlan-id | vlan-name);
        server-reject-vlan (vlan-id | vlan-name);
        server-timeout seconds;
        supplicant (multiple | single | single-secure);
        supplicant-timeout seconds;
        transmit-period seconds;
      }
      static mac-address {
        interface interface-name;
        vlan-assignment (vlan-id |vlan-name);
      }
  }
  gvrp {
```

```
    <enable | disable>;
    interface (all | [interface-name]) {
      disable;
    }
    join-timer millseconds;
    leave-timer milliseconds;
    leaveall-timer milliseconds;
  }
  igmp-snooping {
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>
        <match regex>;
      flag flag (detail | disable | receive | send);
    }
    vlan (vlan-id | vlan-number) {
      data-forwarding {
        source {
          groups group-prefix;
        }
        receiver {
          source-vlans vlan-list;
          install ;
        }
      }
      disable {
        interface interface-name
      }
      immediate-leave;
      interface interface-name {
        group-limit limit;
        multicast-router-interface;
        static {
          group ip-address;
        }
      }
      proxy ;
      query-interval seconds;
      query-last-member-interval seconds;
      query-response-interval seconds;
      robust-count number;
    }
  }
  lldp {
    disable;
    advertisement-interval seconds;
    hold-multiplier number;
    interface (all | interface-name) {
      disable;
    }
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>
        <match regex>;
      flag flag (detail | disable | receive | send);
    }
  }
  lldp-med {
```

```
    disable;
    fast-start number;
    interface (all | interface-name) {
      disable;
      location {
        elin number;
        civic-based {
          what number;
          country-code code;
          ca-type {
            number {
              ca-value value;
            }
          }
        }
      }
    }
  }
  mpls {
    interface ( all | interface-name );
    label-switched-path lsp-name to remote-provider-edge-switch;
    path destination {
      <address | hostname> <strict | loose>
    }
  mstp {
    disable;
    bpdu-block-on-edge;
    bridge-priority priority;
    configuration-name name;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
      disable;
      bpdu-timeout-action {
        block;
        alarm;
      }
      cost  cost;
      edge;
      mode mode;
      no-root-port;
      priority priority;
    }
    max-age seconds;
    max-hops hops;
    msti msti-id {
      vlan (vlan-id | vlan-name);
      interface interface-name {
        disable;
        cost  cost;
        edge;
        mode mode;
        priority priority;
      }
    }
    revision-level revision-level;
```

```
    traceoptions {
      file filename <files number > <size size> <no-stamp | world-readable |
        no-world-readable>;
      flag flag;
    }
}
mvrp {
    disable
    interface (all | interface-name) {
      disable;
      join-timer milliseconds;
      leave-timer milliseconds;
      leaveall-timer milliseconds;
      registration (forbidden | normal);
    }
    no-dynamic-vlan;
    traceoptions {
      file filename <files number > <size size> <no-stamp | world-readable |
        no-world-readable>;
      flag flag;
    }
}
oam {
    ethernet{
      connectivity-fault-management {
        action-profile profile-name {
          default-actions {
            interface-down;
          }
        }
        linktrace {
          age (30m | 10m | 1m | 30s | 10s);
          path-database-size path-database-size;
        }
        maintenance-domain domain-name {
          level number;
          mip-half-function (none | default |explicit);
          name-format (character-string | none | dns | mac+2oct);
          maintenance-association ma-name {
            continuity-check {
              hold-interval minutes;
              interval (10m | 10s | 1m | 1s| 100ms);
              loss-threshold number;
            }
            mep mep-id {
              auto-discovery;
              direction down;
              interface interface-name;
              remote-mep mep-id {
                action-profile profile-name;
              }
            }
          }
        }
      }
      link-fault-management {
```

```
        action-profile profile-name;
          action {
            syslog;
            link-down;
          }
          event {
            link-adjacency-loss;
            link-event-rate;
            frame-error count;
            frame-period count;
            frame-period-summary count;
            symbol-period count;
          }
          interface interface-name {
            link-discovery (active | passive);
            pdu-interval interval;
            event-thresholds threshold-value;
            remote-loopback;
            event-thresholds {
              frame-errorcount;
              frame-period count;
              frame-period-summary count;
              symbol-period count;
            }
          }
          negotiation-options {
            allow-remote-loopback;
            no-allow-link-events;
          }
      }
    }
  }
  rstp {
    disable;
    bpdu-block-on-edge;
    bridge-priority priority;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
      disable;
      bpdu-timeout-action {
        block;
        alarm;
      }
      cost  cost;
      edge;
      mode mode;
      no-root-port;
      priority priority;
    }
    max-age seconds;
    }
    traceoptions {
      file filename <files number > <size size> <no-stamp | world-readable |
        no-world-readable>;
      flag flag;
```

```
    }
  }
  sflow {
    agent-id
    collector {
      ip-address;
      udp-port port-number;
    }
    disable;
    interfaces interface-name {
      disable;
      polling-interval seconds;
      sample-rate number;
    }
    polling-interval seconds;
    sample-rate number;
    source-ip
  }
  stp {
    disable;
    bridge-priority priority;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
      disable;
      bpdu-timeout-action {
        block;
        alarm;
      }
      cost  cost;
      edge;
      mode mode;
      no-root-port;
      priority priority;
    }
    max-age seconds;
  }
    traceoptions {
      file filename <files number > <size size> <no-stamp | world-readable |
        no-world-readable>;
      flag flag;
    }
  vstp {
    bpdu-block-on-edge;
    disable;
    force-version stp;
    vlan (all | vlan-id | vlan-name) {
      bridge-priority priority;
      forward-delay seconds;
      hello-time seconds;
      interface (all | interface-name) {
        bpdu-timeout-action {
          alarm;
          block;
        }
        cost  cost;
```

```
            disable;
            edge;
            mode mode;
            no-root-port;
            priority priority;
          }
          max-age seconds;
          traceoptions {
            file filename <files number > <size size> <no-stamp | world-readable |
              no-world-readable>;
            flag flag;
          }
        }
      }
    }
```

Related Documentation

- 802.1X for J-EX Series Switches Overview on page 2253
- Example: Configure Automatic VLAN Administration Using GVRP on page 1087
- Understanding MAC RADIUS Authentication on J-EX Series Switches
- Understanding Server Fail Fallback and 802.1X Authentication on J-EX Series Switches on page 2258
- IGMP Snooping on J-EX Series Switches Overview on page 2047
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261
- Understanding MSTP for J-EX Series Switches on page 1277
- Understanding Multiple VLAN Registration Protocol (MVRP) on J-EX Series Switches on page 1054
- Understanding Ethernet OAM Connectivity Fault Management for a J-EX Series Switch on page 3463
- Understanding Ethernet OAM Link Fault Management for a J-EX Series Switch on page 3427
- Understanding RSTP for J-EX Series Switches on page 1276
- Understanding STP for J-EX Series Switches on page 1275
- Understanding How to Use sFlow Technology for Network Monitoring on a J-EX Series Switch on page 3283
- Understanding VSTP for J-EX Series Switches on page 1281

alarm

Syntax alarm;

Hierarchy Level [edit protocols **mstp interface** (all | *interface-name*) **bpdu-timeout-action**],
[edit protocols **rstp interface** (all | *interface-name*) **bpdu-timeout-action**],
[edit protocols **stp interface** (all | *interface-name*) **bpdu-timeout-action**],
[edit protocols **vstp vlan** *vlan-id* **interface** (all | *interface-name*) **bpdu-timeout-action**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For interfaces configured for loop protection, configure the software to generate a message to be sent to the system log file to record the loop-protection event.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show spanning-tree bridge on page 1398**
- **show spanning-tree interface on page 1407**
- Example: Configuring Network Regions for VLANs with MSTP on J-EX Series Switches on page 1297
- Example: Configuring Faster Convergence and Improving Network Stability with RSTP on J-EX Series Switches on page 1283
- Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on J-EX Series Switches on page 1325
- Understanding Loop Protection for STP, RSTP, VSTP, and MSTP on J-EX Series Switches on page 1279
- Understanding VSTP for J-EX Series Switches on page 1281

block

Syntax block;

Hierarchy Level [edit protocols **mstp interface** (all | *interface-name*) **bpdu-timeout-action**],
[edit protocols **rstp interface** (all | *interface-name*) **bpdu-timeout-action**],
[edit protocols **stp interface** (all | *interface-name*) **bpdu-timeout-action**],
[edit protocols **vstp vlan** *vlan-id* **interface** (all | *interface-name*) **bpdu-timeout-action**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure loop protection on a specific interface.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show spanning-tree bridge on page 1398**
- **show spanning-tree interface on page 1407**
- Example: Configuring Network Regions for VLANs with MSTP on J-EX Series Switches on page 1297
- Example: Configuring Faster Convergence and Improving Network Stability with RSTP on J-EX Series Switches on page 1283
- Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on J-EX Series Switches on page 1325
- Understanding Loop Protection for STP, RSTP, VSTP, and MSTP on J-EX Series Switches on page 1279
- Understanding VSTP for J-EX Series Switches on page 1281

bpdu-block

Syntax

```
bpdu-block {
  interface (all | [interface-name]);
  disable-timeout timeout;
}
```

Hierarchy Level [edit ethernet-switching-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure BPDU protection on an interface. If the interface receives BPDUs, it is disabled.

The statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show spanning-tree bridge on page 1398**
- **show spanning-tree interface on page 1407**
- **clear ethernet-switching bpdu-error on page 1390**
- Example: Configuring BPDU Protection on non-STP Interfaces to Prevent STP Miscalculations on J-EX Series Switches on page 1321
- Unblocking an Interface That Receives BPDUs in Error (CLI Procedure) on page 1335
- Example: Configuring Network Regions for VLANs with MSTP on J-EX Series Switches on page 1297
- Example: Configuring Faster Convergence and Improving Network Stability with RSTP on J-EX Series Switches on page 1283

bpdu-block-on-edge

Syntax bpdu-block-on-edge;

Hierarchy Level [edit protocols **mstp**],
[edit protocols **rstp**],
[edit protocols **vstp**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure bridge protocol data unit (BPDU) protection on all edge ports of a switch.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show spanning-tree bridge on page 1398**
- **show spanning-tree interface on page 1407**
- **clear ethernet-switching bpdu-error on page 1390**
- Example: Configuring Network Regions for VLANs with MSTP on J-EX Series Switches on page 1297
- Example: Configuring Faster Convergence and Improving Network Stability with RSTP on J-EX Series Switches on page 1283
- Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations on J-EX Series Switches on page 1317
- Understanding VSTP for J-EX Series Switches on page 1281

bpdu-timeout-action

Syntax

```
bpdu-timeout-action {
  block;
  alarm;
}
```

Hierarchy Level

[edit protocols **mstp interface** (all | *interface-name*)],
[edit protocols **rstp interface** (all | *interface-name*)],
[edit protocols **stp interface** (all | *interface-name*)],
[edit protocols **vstp vlan** *vlan-id* **interface** (all | *interface-name*)]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure the BPDU timeout action on a specific interface. You must configure at least one action (**alarm**, **block**, or both).

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show spanning-tree bridge on page 1398**
- **show spanning-tree interface on page 1407**
- Example: Configuring Network Regions for VLANs with MSTP on J-EX Series Switches on page 1297
- Example: Configuring Faster Convergence and Improving Network Stability with RSTP on J-EX Series Switches on page 1283
- Example: Configuring Loop Protection to Prevent Interfaces from Transitioning from Blocking to Forwarding in a Spanning Tree on J-EX Series Switches on page 1325
- Understanding Loop Protection for STP, RSTP, VSTP, and MSTP on J-EX Series Switches on page 1279
- Understanding VSTP for J-EX Series Switches on page 1281

bridge-priority

Syntax bridge-priority *priority*;

Hierarchy Level [edit protocols **mstp**],
[edit protocols **mstp msti** *msti-id*],
[edit protocols **rstp**],
[edit protocols **stp**],
[edit protocols **vstp vlan** *vlan-id*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the bridge priority. The bridge priority determines which bridge is elected as the root bridge. If two bridges have the same path cost to the root bridge, the bridge priority determines which bridge becomes the designated bridge for a LAN segment.

Default 32,768

Options ***priority***—Bridge priority. It can be set only in increments of 4096.

Range: 0 through 61,440

Default: 32,768

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show spanning-tree bridge on page 1398**
- **show spanning-tree interface on page 1407**
- Example: Configuring Network Regions for VLANs with MSTP on J-EX Series Switches on page 1297
- Understanding MSTP for J-EX Series Switches on page 1277
- Understanding VSTP for J-EX Series Switches on page 1281

configuration-name

Syntax configuration-name *configuration-name*;

Hierarchy Level [edit protocols **mstp**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the configuration name. The configuration name is the MSTP region name carried in the MSTP BPDUs.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show spanning-tree bridge on page 1398**
- **show spanning-tree interface on page 1407**
- Example: Configuring Network Regions for VLANs with MSTP on J-EX Series Switches on page 1297
- Example: Configuring Faster Convergence and Improving Network Stability with RSTP on J-EX Series Switches on page 1283
- Understanding MSTP for J-EX Series Switches on page 1277

cost

Syntax cost *cost*;

Hierarchy Level [edit protocols **mstp interface** (all | *interface-name*)],
[edit protocols **mstp msti** *msti-id* **interface** *interface-name*],
[edit protocols **rstp interface** (all | *interface-name*)],
[edit protocols **stp interface** (all | *interface-name*)],
[edit protocols **vstp vlan** *vlan-id* **interface** (all | *interface-name*)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), configure the link cost to control which bridge is the designated bridge and which interface is the designated interface.

Default The link cost is determined by the link speed.

Options ***cost***—Link cost associated with the port.

Range: 1 through 200,000,000

Default: Link cost is determined by the link speed.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show spanning-tree bridge on page 1398**
- **show spanning-tree interface on page 1407**
- Understanding STP for J-EX Series Switches on page 1275
- Understanding MSTP for J-EX Series Switches on page 1277
- Understanding VSTP for J-EX Series Switches on page 1281

disable

Syntax disable;

Hierarchy Level
[edit protocols **mstp**],
[edit protocols **mstp interface** *interface-name*],
[edit protocols **mstp msti** *msti-id* **vlan** (*vlan-id* | *vlan-name*) **interface** *interface-name*],
[edit protocols **rstp**],
[edit protocols **rstp interface** *interface-name*],
[edit protocols **stp**],
[edit protocols **stp interface** *interface-name*],
[edit protocols **vstp**],
[edit protocols **vstp vlan** *vlan-id* **interface** (all | *interface-name*)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable STP, MSTP, RSTP, or VSTP on the switch or on a specific interface.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show spanning-tree bridge on page 1398**
- **show spanning-tree interface on page 1407**
- Example: Configuring Network Regions for VLANs with MSTP on J-EX Series Switches on page 1297
- Example: Configuring Faster Convergence and Improving Network Stability with RSTP on J-EX Series Switches on page 1283
- Understanding MSTP for J-EX Series Switches on page 1277
- Understanding STP for J-EX Series Switches on page 1275
- Understanding VSTP for J-EX Series Switches on page 1281

disable-timeout

Syntax disable-timeout *timeout*;

Hierarchy Level [edit **ethernet-switching-options bpdu-block**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For interfaces configured for BPDU protection, specify the amount of time an interface receiving BPDUs is disabled.

Default The disable timeout is not enabled.

Options ***timeout*** —Amount of time, in seconds, the interface receiving BPDUs is disabled. Once the timeout expires, the interface is brought back into service.

Range: 10 through 3600 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show spanning-tree bridge on page 1398**
- **show spanning-tree interface on page 1407**
- Example: Configuring Network Regions for VLANs with MSTP on J-EX Series Switches on page 1297
- Example: Configuring Faster Convergence and Improving Network Stability with RSTP on J-EX Series Switches on page 1283
- Example: Configuring BPDU Protection on non-STP Interfaces to Prevent STP Miscalculations on J-EX Series Switches on page 1321
- Understanding BPDU Protection for STP, RSTP, and MSTP on J-EX Series Switches on page 1278

edge

Syntax edge;

Hierarchy Level [edit protocols **mstp interface** (all | *interface-name*)],
[edit protocols **mstp msti** *msti-id* **interface** *interface-name*],
[edit protocols **rstp interface** (all | *interface-name*)],
[edit protocols **stp interface** (all | *interface-name*)],
[edit protocols **vstp vlan** *vlan-id* **interface** (all | *interface-name*)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), configure interfaces as edge interfaces. Edge interfaces immediately transition to a forwarding state.

Default Edge interfaces are not enabled.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show spanning-tree bridge on page 1398**
- **show spanning-tree interface on page 1407**
- Example: Configuring Network Regions for VLANs with MSTP on J-EX Series Switches on page 1297
- Example: Configuring Faster Convergence and Improving Network Stability with RSTP on J-EX Series Switches on page 1283
- Understanding MSTP for J-EX Series Switches on page 1277
- Understanding STP for J-EX Series Switches on page 1275
- Understanding VSTP for J-EX Series Switches on page 1281

force-version

Syntax force-version stp;

Hierarchy Level [edit protocols vstp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Force VLAN Spanning Tree Protocol (VSTP) to use the STP protocol instead of the default protocol, RSTP.

Options **stp**—Spanning Tree Protocol

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show spanning-tree bridge on page 1398**
- **show spanning-tree interface on page 1407**
- Understanding VSTP for J-EX Series Switches on page 1281

forward-delay

Syntax forward-delay *seconds*;

Hierarchy Level [edit protocols **mstp**],
[edit protocols **rstp**],
[edit protocols **stp**],
[edit protocols **vstp vlan** *vlan-id*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), specify how long a bridge interface remains in the listening and learning states before transitioning to the forwarding state.

Default 15 seconds

Options ***seconds***—Number of seconds the bridge interface remains in the listening and learning states.

Range: 4 through 30 seconds

Default: 15 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show spanning-tree bridge on page 1398**
- **show spanning-tree interface on page 1407**
- Example: Configuring Network Regions for VLANs with MSTP on J-EX Series Switches on page 1297
- Example: Configuring Faster Convergence and Improving Network Stability with RSTP on J-EX Series Switches on page 1283
- Understanding MSTP for J-EX Series Switches on page 1277
- Understanding STP for J-EX Series Switches on page 1275
- Understanding VSTP for J-EX Series Switches on page 1281

hello-time

Syntax hello-time *seconds*;

Hierarchy Level [edit protocols **mstp**],
[edit protocols **rstp**],
[edit protocols **stp**],
[edit protocols **vstp vlan** *vlan-id*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), specify the time interval at which the root bridge transmits configuration BPDUs.

Default 2 seconds

Options ***seconds***—Number of seconds between transmissions of configuration BPDUs.

Range: 1 through 10 seconds

Default: 2 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show spanning-tree bridge on page 1398**
- **show spanning-tree interface on page 1407**
- Example: Configuring Network Regions for VLANs with MSTP on J-EX Series Switches on page 1297
- Example: Configuring Faster Convergence and Improving Network Stability with RSTP on J-EX Series Switches on page 1283
- Understanding MSTP for J-EX Series Switches on page 1277
- Understanding STP for J-EX Series Switches on page 1275
- Understanding VSTP for J-EX Series Switches on page 1281

interface

Syntax interface (all | [*interface-name*]);

Hierarchy Level [edit **ethernet-switching-options bpdu-block**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply BPDU protection to all interfaces or one or more interfaces.

Options **all**—All interfaces.

interface-name —Name of a Gigabit Ethernet interface.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show spanning-tree bridge on page 1398**
- **show spanning-tree interface on page 1407**
- Example: Configuring Network Regions for VLANs with MSTP on J-EX Series Switches on page 1297
- Example: Configuring Faster Convergence and Improving Network Stability with RSTP on J-EX Series Switches on page 1283
- Example: Configuring BPDU Protection on non-STP Interfaces to Prevent STP Miscalculations on J-EX Series Switches on page 1321
- Understanding BPDU Protection for STP, RSTP, and MSTP on J-EX Series Switches on page 1278

interface

Syntax

```
interface interface-name {
    disable;
    cost cost;
    edge;
    mode mode;
    no-root-port;
    priority priority;
}
```

Hierarchy Level

[edit protocols **mstp**],
[edit protocols **mstp msti** *msti-id*],
[edit protocols **rstp**],
[edit protocols **stp**],
[edit protocols **vstp vlan** *vlan-id*]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), configure an interface.

The **edge**, **mode**, and **no-root-port** options are not available at the [**edit protocols mstp msti** ***msti-id***] hierarchy level.

Options

interface-name—Name of a Gigabit Ethernet interface.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show spanning-tree bridge on page 1398**
- **show spanning-tree interface on page 1407**
- Example: Configuring Network Regions for VLANs with MSTP on J-EX Series Switches on page 1297
- Example: Configuring Faster Convergence and Improving Network Stability with RSTP on J-EX Series Switches on page 1283
- Understanding MSTP for J-EX Series Switches on page 1277
- Understanding RSTP for J-EX Series Switches on page 1276
- Understanding STP for J-EX Series Switches on page 1275
- Understanding VSTP for J-EX Series Switches on page 1281

max-age

Syntax max-age *seconds*;

Hierarchy Level [edit protocols **mstp**],
[edit protocols **rstp**],
[edit protocols **stp**],
[edit protocols **vstp vlan** *vlan-id*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), specify the maximum age of received protocol BPDUs.

Default 20 seconds

Options ***seconds***—The maximum age of received protocol BPDUs.

Range: 6 through 40 seconds

Default: 20 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show spanning-tree bridge on page 1398**
- **show spanning-tree interface on page 1407**
- Example: Configuring Network Regions for VLANs with MSTP on J-EX Series Switches on page 1297
- Example: Configuring Faster Convergence and Improving Network Stability with RSTP on J-EX Series Switches on page 1283
- Understanding MSTP for J-EX Series Switches on page 1277
- Understanding STP for J-EX Series Switches on page 1275
- Understanding VSTP for J-EX Series Switches on page 1281

max-hops

Syntax max-hops *hops*;

Hierarchy Level [edit protocols mstp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For Multiple Spanning Tree Protocol (MSTP), configure the maximum number of hops a BPDU can be forwarded in the MSTP region.

Default 20 hops

Options ***hops*** — Number of hops the BPDU can be forwarded.
Range: 1 through 255 hops
Default: 20 hops

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show spanning-tree bridge on page 1398**
- **show spanning-tree interface on page 1407**
- Example: Configuring Network Regions for VLANs with MSTP on J-EX Series Switches on page 1297
- Understanding MSTP for J-EX Series Switches on page 1277

mode

Syntax mode *mode*;

Hierarchy Level [edit protocols **mstp interface** (all | *interface-name*)],
[edit protocols **mstp msti** *msti-id* **interface** *interface-name*],
[edit protocols **rstp interface** (all | *interface-name*)],
[edit protocols **stp interface** (all | *interface-name*)],
[edit protocols **vstp vlan** *vlan-id* **interface** (all | *interface-name*)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), configure the link mode to identify point-to-point links.

Default For a full-duplex link, the default link mode is **point-to-point**. For a half-duplex link, the default link mode is **shared**.

Options *mode*—Link mode:

- **point-to-point**—Link is point to point.
- **shared**—Link is shared media.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show spanning-tree bridge on page 1398**
- **show spanning-tree interface on page 1407**
- Example: Configuring Network Regions for VLANs with MSTP on J-EX Series Switches on page 1297
- Example: Configuring Faster Convergence and Improving Network Stability with RSTP on J-EX Series Switches on page 1283
- Understanding MSTP for J-EX Series Switches on page 1277
- Understanding STP for J-EX Series Switches on page 1275
- Understanding VSTP for J-EX Series Switches on page 1281

msti

Syntax

```
msti msti-id {
  vlan (vlan-id | vlan-name);
  interface interface-name {
    disable;
    cost cost;
    priority priority;
  }
}
```

Hierarchy Level [edit protocols mstp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the Multiple Spanning Tree Instance (MSTI) identifier for Multiple Spanning Tree Protocol (MSTP). MSTI IDs are local to each region, so you can reuse the same MSTI ID in different regions.

Default MSTI is disabled.

Options *msti-id* —MSTI identifier.

Range: 1 through 4094. The Common Instance Spanning Tree (CIST) is always MSTI 0.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show spanning-tree bridge on page 1398**
- **show spanning-tree interface on page 1407**
- Example: Configuring Network Regions for VLANs with MSTP on J-EX Series Switches on page 1297
- Understanding MSTP for J-EX Series Switches on page 1277

mstp

Syntax

```
mstp {
  disable;
  bpdu-block-on-edge;
  bridge-priority priority;
  configuration-name name;
  forward-delay seconds;
  hello-time seconds;
  interface ( all | interface-name {
    bpdu-timeout-action {
      block;
      alarm;
    }
    disable;
    cost cost;
    edge;
    mode mode;
    no-root-port;
    priority priority;
  }
  max-age seconds;
  max-hops hops;
  msti msti-id {
    vlan (vlan-id | vlan-name);
    interface interface-name {
      disable;
      cost cost;
      priority priority;
    }
  }
  traceoptions {
    file filename <files number > <size size> <no-stamp | world-readable |
      no-world-readable>;
    flag flag;
  }
  revision-level revision-level;
}
```

Hierarchy Level [edit protocols]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure Multiple Spanning Tree Protocol (MSTP). MSTP is defined in the IEEE 802.1Q-2003 specification and is used to create a loop-free topology in networks with multiple spanning tree regions.

The statements are explained separately.

Default MSTP is disabled.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show spanning-tree bridge on page 1398**
- **show spanning-tree interface on page 1407**
- Example: Configuring Network Regions for VLANs with MSTP on J-EX Series Switches on page 1297
- Understanding MSTP for J-EX Series Switches on page 1277

no-root-port

Syntax

no-root-port;

Hierarchy Level

[edit protocols **mstp interface** (all | *interface-name*)],
[edit protocols **rstp interface** (all | *interface-name*)],
[edit protocols **stp interface** (all | *interface-name*)],
[edit protocols **vstp vlan** *vlan-id* **interface** (all | *interface-name*)]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure an interface to be a spanning tree designated port. If the bridge receives superior STP bridge protocol data units (BPDUs) on a root-protected interface, that interface transitions to a root-prevented STP state (inconsistency state) and the interface is blocked. This blocking prevents a bridge that should not be the root bridge from being elected the root bridge. When the bridge stops receiving superior STP BPDUs on the root-protected interface, interface traffic is no longer blocked.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show spanning-tree bridge on page 1398**
- **show spanning-tree interface on page 1407**
- Example: Configuring Root Protection to Enforce Root Bridge Placement in Spanning Trees on J-EX Series Switches on page 1329
- Example: Configuring Network Regions for VLANs with MSTP on J-EX Series Switches on page 1297
- Example: Configuring Faster Convergence and Improving Network Stability with RSTP on J-EX Series Switches on page 1283
- Understanding VSTP for J-EX Series Switches on page 1281

priority

Syntax priority *priority*;

Hierarchy Level [edit protocols **mstp interface** (all | *interface-name*)],
[edit protocols **mstp msti** *msti-id* **interface** *interface-name*],
[edit protocols **rstp interface** (all | *interface-name*)],
[edit protocols **stp interface** (all | *interface-name*)],
[edit protocols **vstp vlan** *vlan-id* **interface** (all | *interface-name*)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), VLAN Spanning Tree Protocol (VSTP), or Multiple Spanning Tree Protocol (MSTP), specify the interface priority to control which interface is elected as the root port.

Default The default value is 128.

Options ***priority***—Interface priority. The interface priority must be set in increments of 16.
Range: 0 through 240

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show spanning-tree bridge on page 1398**
- **show spanning-tree interface on page 1407**
- Example: Configuring Network Regions for VLANs with MSTP on J-EX Series Switches on page 1297
- Example: Configuring Faster Convergence and Improving Network Stability with RSTP on J-EX Series Switches on page 1283
- Understanding MSTP for J-EX Series Switches on page 1277
- Understanding STP for J-EX Series Switches on page 1275
- Understanding VSTP for J-EX Series Switches on page 1281

revision-level

Syntax revision-level *revision-level*;

Hierarchy Level [edit protocols mstp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For Multiple Spanning Tree Protocol (MSTP), set the revision number of the MSTP configuration.

Default The revision level is disabled.

Options ***revision-level*** —Revision number of the MSTP region configuration.

Range: 0 through 65535

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show spanning-tree bridge on page 1398**
- **show spanning-tree interface on page 1407**
- Example: Configuring Network Regions for VLANs with MSTP on J-EX Series Switches on page 1297
- Understanding MSTP for J-EX Series Switches on page 1277

rstp

Syntax

```
rstp {
  disable;
  bpdu-block-on-edge;
  bridge-priority priority;
  forward-delay seconds;
  hello-time seconds;
  interface (all | interface-name) {
    disable;
    bpdu-timeout-action{
      alarm;
      block;
    }
    cost cost;
    edge;
    mode mode;
    no-root-port;
    priority priority;
  }
  max-age seconds;
  traceoptions {
    file filename <files number > <size size> <no-stamp | no-world-readable |
      world-readable>;
    flag flag;
  }
}
```

Hierarchy Level [edit protocols]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure Rapid Spanning Tree Protocol (RSTP). RSTP is defined in the IEEE 802.1D-2004 specification and is used to prevent loops in Layer 2 networks, which results in shorter convergence times than those provided by basic Spanning Tree Protocol (STP).

VSTP and RSTP can be configured concurrently. You can selectively configure up to 253 VLANs using VSTP; the remaining VLANs will be configured using RSTP. VSTP and RSTP are the only spanning-tree protocols that can be configured concurrently on the switch. See Configuring VSTP (CLI Procedure) for more information on configuring VSTP and RSTP concurrently.

BEST PRACTICE: Configure RSTP when you configure VSTP. RSTP overhead is minimal and this configuration ensures that a spanning-tree protocol is running on all VLANs on your switch, even when your switch is supporting more than 253 VLANs.

The remaining statements are explained separately.

Default RSTP is enabled on all Ethernet switching interfaces.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show spanning-tree bridge on page 1398**
- **show spanning-tree interface on page 1407**
- Example: Configuring Faster Convergence and Improving Network Stability with RSTP on J-EX Series Switches on page 1283
- Understanding RSTP for J-EX Series Switches on page 1276

stp

Syntax

```
stp {
  disable;
  bridge-priority priority;
  forward-delay seconds;
  hello-time seconds;
  interface (all | interface-name) {
    disable;
    bpdu-timeout-action {
      block;
      alarm;
    }
    cost  cost;
    edge;
    mode mode;
    no-root-port;
    priority priority;
  }
  max-age seconds;
  traceoptions {
    file filename <files number > <size size> <no-stamp | world-readable |
      no-world-readable>;
    flag flag;
  }
}
```

Hierarchy Level [edit protocols]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description When you explicitly configure STP, the J-EX Series switches use the IEEE 802.1D 2004 specification, force version 0. This configuration runs a version of RSTP that is compatible with the classic, basic STP (defined in the IEEE 802.1D 1998 specification).

The remaining statements are explained separately.

Default STP is disabled; by default, RSTP is enabled on all Ethernet switching ports.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show spanning-tree bridge on page 1398**
- **show spanning-tree interface on page 1407**
- Example: Configuring BPDU Protection on STP Interfaces to Prevent STP Miscalculations on J-EX Series Switches on page 1317
- Configuring STP (CLI Procedure) on page 1336
- Understanding STP for J-EX Series Switches on page 1275

traceoptions

Syntax

```
traceoptions {
  file name <replace> <size size> <files number> <no-stamp>
    <(world-readable | no-world-readable)>;
  flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

[edit protocols **mstp**],
[edit protocols **rstp**],
[edit protocols **stp**],
[edit protocols **vstp vlan** *vlan-id*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set protocol-level tracing options for STP, RSTP, MSTP, and VSTP.

Default Traceoptions is disabled.

Options

disable—(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file ***name*** —Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. We recommend that you place STP tracing output in the file **/var/log/stp-log**.

files ***number*** —(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*** **.0**, then ***trace-file*** **.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum number of files, you must also specify a maximum file size with the size option.

Range: 2 through 1000 files

Default: 1 trace file only

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements:

- **all** —Trace all operations.
- **all-failures**—Trace all failure conditions.
- **bpdu** —Trace BPDU reception and transmission. Note that you must also use port-transmit-state-machine in order to log transmit operations.
- **bridge-detection-state-machine** —Trace the bridge detection state machine.
- **events** —Trace events of the protocol state machine.

- **port-information-state-machine** —Trace the port information state machine.
- **port-migration-state-machine** —Trace the port migration state machine.
- **port-receive-state-machine** —Trace the port receive state machine.
- **port-role-select-state-machine** —Trace the port role selection state machine.
- **port-role-transit-state-machine** —Trace the port role transit state machine.
- **port-state-transit-state-machine** —Trace the port state transit state machine.
- **port-transmit-state-machine** —Trace the port transmit state machine
- **ppmd** —Trace the state and events for the ppmd process
- **state-machine-variables** —Trace when the state machine variables change
- **timers** —Trace protocol timers
- **topology-change-state-machine** —Trace the topology change state machine.

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

Default: If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Prevent aney user from reading the log file.

replace—(Optional) Replace an existing trace file if there is one.

Default: If you do not include this option, tracing output is appended to an existing trace file.

size ***size*** —(Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*** **.0.** When the ***trace-file*** again reaches its maximum size, ***trace-file*** **.0** is renamed ***trace-file*** **.1** and ***trace-file*** is renamed ***trace-file*** **.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show spanning-tree bridge on page 1398**
- **show spanning-tree interface on page 1407**
- Example: Configuring Network Regions for VLANs with MSTP on J-EX Series Switches on page 1297
- Example: Configuring Faster Convergence and Improving Network Stability with RSTP on J-EX Series Switches on page 1283
- Understanding MSTP for J-EX Series Switches on page 1277
- Understanding RSTP for J-EX Series Switches on page 1276
- Understanding STP for J-EX Series Switches on page 1275
- Understanding VSTP for J-EX Series Switches on page 1281

vlan

Syntax

```
vlan (vlan-id | vlan-name) {
  bridge-priority priority;
  forward-delay seconds;
  hello-time seconds;
  interface interface-name {
    bpdu-timeout-action {
      alarm;
      block;
    }
    cost  cost;
    disable;
    edge;
    mode mode;
    no-root-port;
    priority priority;
  }
  max-age seconds;
  traceoptions {
    file filename <files number > <size size> <no-stamp | world-readable |
      no-world-readable>;
    flag flag;
  }
}
```

Hierarchy Level [edit protocols **mstp msti** *msti-id*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the VLANs for a Multiple Spanning Tree Instance (MSTI).

TIP: To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type ? after vlan or vlans in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.

Default Not enabled.

Options ***vlan-id***—Numeric VLAN identifier.

vlan-name—Name of the VLAN.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Network Regions for VLANs with MSTP on J-EX Series Switches on page 1297

- Understanding MSTP for J-EX Series Switches on page 1277

vlan (VSTP)

Syntax

```
vlan (all | vlan-id | vlan-name) {
  bridge-priority priority;
  forward-delay seconds;
  hello-time seconds;
  interface (all | interface-name) {
    bpdu-timeout-action {
      alarm;
      block;
    }
    cost  cost;
    disable;
    edge;
    mode mode;
    no-root-port;
    priority priority;
  }
  max-age seconds;
  traceoptions {
    file filename <files number > <size size> <no-stamp | world-readable |
      no-world-readable>;
    flag flag;
  }
}
```

Hierarchy Level [edit protocols vstp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure VSTP VLAN parameters.

TIP: To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type ? after vlan or vlans in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.

Options **all**—All VLANs.

vlan-id—Numeric VLAN identifier.

vlan-name—Name of the VLAN.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Understanding VSTP for J-EX Series Switches on page 1281

vstp

Syntax

```
vstp {
  bpdu-block-on-edge;
  disable;
  force-version stp;
  vlan (vlan-id | vlan-name) {
    bridge-priority priority;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
      disable;
      bpdu-timeout-action{
        alarm;
        block;
      }
      cost  cost;
      edge;
      mode mode;
      no-root-port;
      priority priority;
    }
    max-age seconds;
    traceoptions {
      file filename <files number > <size size> <no-stamp | no-world-readable |
        world-readable>;
      flag flag;
    }
  }
}
```

Hierarchy Level [edit protocols]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure VLAN Spanning Tree Protocol (VSTP). VSTP is used to prevent loops in Layer 2 networks on a per-VLAN basis.

You cannot commit a configuration that uses VSTP on more than 253 VLANs. If there are more than 253 VLANs on your switch, you must use the **vlan** statement to specify which VLANs or VLAN groups should use VSTP, and the total number of VLANs using VSTP cannot exceed 253. You also cannot use the **vlan all** option to configure VSTP when your switch has more than 253 VLANs. Run RSTP with VSTP in networks with large numbers of VLANs to ensure all VLANs are running a spanning-tree protocol.

BEST PRACTICE: Configure RSTP when you configure VSTP. RSTP overhead is minimal and this configuration ensures that a spanning-tree protocol is running on all VLANs on your switch, even when your switch is supporting more than 253 VLANs.

The remaining statements are explained separately.

Default VSTP is not enabled by default.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show spanning-tree bridge on page 1398**
- **show spanning-tree interface on page 1407**
- Configuring VLAN Spanning Tree Protocol (CLI Procedure) on page 1340
- Understanding VSTP for J-EX Series Switches on page 1281

CHAPTER 69

Operational Mode Commands for Spanning-Tree Protocols

clear ethernet-switching bpdu-error

Syntax clear ethernet-switching bpdu-error interface *interface-name*

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Clear bridge protocol data unit (BPDU) errors from an interface and unblock the interface.

Options ***interface-name*** —Clear BPDU errors on the specified interface.

Required Privilege Level clear

Related Documentation

- **show spanning-tree statistics on page 1416**
- Understanding BPDU Protection for STP, RSTP, and MSTP on J-EX Series Switches on page 1278

List of Sample Output **clear ethernet-switching bpdu-error interface ge-0/0/1.0 on page 1390**

clear ethernet-switching bpdu-error interface ge-0/0/1.0

```
user@switch> clear ethernet-switching bpdu-error interface ge-0/0/1.0
```

clear spanning-tree statistics

Syntax
clear spanning-tree statistics
<interface *interface-name*>
<logical-system *logical-system-name*>

Syntax (J-EX Series Switch)
clear spanning-tree statistics
<interface *interface-name*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Clear Spanning Tree Protocol statistics.

Options
none—Reset STP counters for all interfaces for all routing instances.

interface *interface-name*—(Optional) Clear STP statistics for the specified interface only.

logical-system *logical-system-name*—(Optional) Clear STP statistics on a particular logical system.

Required Privilege Level
clear

Related Documentation
- **show spanning-tree statistics on page 1414**

List of Sample Output
clear stp statistics on page 1391

clear stp statistics
```
user@host> clear stp statistics
```

clear spanning-tree statistics

Syntax
clear spanning-tree statistics
<interface *interface-name* unit *logical-unit-number*>;

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Reset STP statistics for the all interfaces or a specified interface.

Options
none—Reset STP counters for all interfaces.

interface-name —(Optional) The name of the interface for which statistics should be reset.

logical-unit-number —(Optional) The logical unit number of the interface.

Required Privilege Level
clear

Related Documentation
- **show spanning-tree bridge on page 1398**
- **show spanning-tree interface on page 1407**
- Understanding STP for J-EX Series Switches on page 1275

List of Sample Output
clear spanning-tree statistics on page 1392

Output Fields
This command produces no output.

clear spanning-tree statistics
```
user@switch> clear spanning—tree statistics
```

show spanning-tree bridge

Syntax

show spanning-tree bridge
<brief | detail>
<msti *msti-id*>
<routing-instance *routing-instance-name*>
<vlan-id *vlan-id*>

Syntax (J-EX Series Switch)

show spanning-tree bridge
<brief | detail>
<msti *msti-id*>
<vlan-id *vlan-id*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display the configured or calculated Spanning Tree Protocol (STP) parameters.

Options

none—(Optional) Display brief STP bridge information for all multiple spanning-tree instances (MSTIs).

brief | detail—(Optional) Display the specified level of output.

msti *msti-id*—(Optional) Display STP bridge information for the specified MSTI.

routing-instance *routing-instance-name*—(Optional) Display STP bridge information for the specified routing instance.

vlan-id *vlan-id*—(Optional) Display STP bridge information for the specified VLAN.

Required Privilege Level

view

List of Sample Output

show spanning-tree bridge routing-instance on page 1394
show spanning-tree bridge msti on page 1395
show spanning-tree bridge vlan-id (MSTP) on page 1395
show spanning-tree bridge (VSTP) on page 1396
show spanning-tree bridge vlan-id (VSTP) on page 1397

Output Fields

Table 179 on page 1393 lists the output fields for the **show spanning-tree bridge** command. Output fields are listed in the approximate order in which they appear.

Table 179: show spanning-tree bridge Output Fields

Field Name	Field Description
Routing instance name	Name of the routing instance under which the bridge is configured.
Enabled protocol	Spanning Tree Protocol type enabled.
Root ID	Bridge ID of the elected spanning tree root bridge. The bridge ID consists of a configurable bridge priority and the MAC address of the bridge.

Table 179: show spanning-tree bridge Output Fields *(continued)*

Field Name	Field Description
Root cost	Calculated cost to reach the root bridge from the bridge where the command is entered.
Root port	Interface that is the current elected root port for this bridge.
CIST regional root	Bridge ID of the elected MSTP regional root bridge.
CIST internal root cost	Calculated cost to reach the regional root bridge from the bridge where the command is entered.
Hello time	Configured number of seconds between transmissions of configuration bridge protocol data units (BPDUs).
Maximum age	Configured maximum expected arrival time of hello bridge protocol data units (BPDUs).
Forward delay	Configured time an STP bridge port remains in the listening and learning states before transitioning to the forwarding state.
Hop count	Configured maximum number of hops a BPDU can be forwarded in the MSTP region.
Message age	Number of elapsed seconds since the most recent BPDU was received.
Number of topology changes	Total number of STP topology changes detected since the routing device last booted.
Time since last topology change	Number of elapsed seconds since the most recent topology change.
Bridge ID (Local)	Locally configured bridge ID. The bridge ID consists of a configurable bridge priority and the MAC address of the bridge.
Extended system ID	System identifier.
MSTI regional root	Bridge ID of the elected MSTP regional root bridge.

show spanning-tree bridge routing-instance

```
user@host> show spanning-tree bridge routing-instance vs1 detail
STP bridge parameters
Routing instance name              : vs1
Enabled protocol                   : MSTP

STP bridge parameters for CIST
  Root ID                          : 32768.00:13:c3:9e:c8:80
  Root cost                        : 0
  Root port                        : xe-10/2/0
  CIST regional root               : 32768.00:13:c3:9e:c8:80
  CIST internal root cost          : 22000
  Hello time                       : 2 seconds
```

```
  Maximum age                        : 20 seconds
  Forward delay                      : 15 seconds
  Hop count                          : 18
  Message age                        : 0
  Number of topology changes         : 1
  Time since last topology change    : 1191 seconds
  Local parameters
    Bridge ID                        : 32768.00:90:69:0b:7f:d1
    Extended system ID               : 1

STP bridge parameters for MSTI 1
  MSTI regional root                 : 32769.00:13:c3:9e:c8:80
  Root cost                          : 22000
  Root port                          : xe-10/2/0
  Hello time                         : 2 seconds
  Maximum age                        : 20 seconds
  Forward delay                      : 15 seconds
  Hop count                          : 18
  Number of topology changes         : 1
  Time since last topology change    : 1191 seconds
  Local parameters
    Bridge ID                        : 32769.00:90:69:0b:7f:d1
    Extended system ID               : 1

STP bridge parameters for MSTI 2
  MSTI regional root                 : 32770.00:13:c3:9e:c8:80
  Root cost                          : 22000
  Root port                          : xe-10/2/0
  Hello time                         : 2 seconds
  Maximum age                        : 20 seconds
  Forward delay                      : 15 seconds
  Hop count                          : 18
  Number of topology changes         : 1
  Time since last topology change    : 1191 seconds
  Local parameters
    Bridge ID                        : 32770.00:90:69:0b:7f:d1
    Extended system ID               : 1
```

show spanning-tree bridge msti

```
user@host> show spanning-tree bridge msti 1 routing-instance vs1 detail
STP bridge parameters
Routing instance name                : vs1
Enabled protocol                     : MSTP

STP bridge parameters for MSTI 1
  MSTI regional root                 : 32769.00:13:c3:9e:c8:80
  Root cost                          : 22000
  Root port                          : xe-10/2/0
  Hello time                         : 2 seconds
  Maximum age                        : 20 seconds
  Forward delay                      : 15 seconds
  Hop count                          : 18
  Number of topology changes         : 1
  Time since last topology change    : 1191 seconds
  Local parameters
    Bridge ID                        : 32769.00:90:69:0b:7f:d1
    Extended system ID               : 1
```

show spanning-tree bridge vlan-id (MSTP)

```
user@host> show spanning-tree bridge vlan-id 1 101 routing-instance vs1 detail
STP bridge parameters
Routing instance name                : vs1
Enabled protocol                     : MSTP
```

```
STP bridge parameters for CIST
  Root ID                           : 32768.00:13:c3:9e:c8:80
  Root cost                         : 0
  Root port                         : xe-10/2/0
  CIST regional root                : 32768.00:13:c3:9e:c8:80
  CIST internal root cost           : 22000
  Hello time                        : 2 seconds
  Maximum age                       : 20 seconds
  Forward delay                     : 15 seconds
  Hop count                         : 18
  Message age                       : 0
  Number of topology changes        : 0
  Local parameters
    Bridge ID                       : 32768.00:90:69:0b:7f:d1
    Extended system ID              : 1
    Hello time                      : 2 seconds
    Maximum age                     : 20 seconds
    Forward delay                   : 15 seconds
    Path cost method                : 32 bit
    Maximum hop count               : 20
```

show spanning-tree bridge (VSTP)

```
user@host> show spanning-tree bridge
STP bridge parameters
Routing instance name               : GLOBAL
Enabled protocol                    : RSTP
  Root ID                           : 28672.00:90:69:0b:3f:d0
  Hello time                        : 2 seconds
  Maximum age                       : 20 seconds
  Forward delay                     : 15 seconds
  Message age                       : 0
  Number of topology changes        : 58
  Time since last topology change   : 14127 seconds
  Local parameters
    Bridge ID                       : 28672.00:90:69:0b:3f:d0
    Extended system ID              : 0

 STP bridge parameters for bridge VLAN 10
   Root ID                          : 28672.00:90:69:0b:3f:d0
   Hello time                       : 2 seconds
   Maximum age                      : 20 seconds
   Forward delay                    : 15 seconds
   Message age                      : 0
   Number of topology changes       : 58
   Time since last topology change  : 14127 seconds
   Local parameters
     Bridge ID                      : 28672.00:90:69:0b:3f:d0
     Extended system ID             : 0

 STP bridge parameters for bridge VLAN 20
   Root ID                          : 28672.00:90:69:0b:3f:d0
   Hello time                       : 2 seconds
   Maximum age                      : 20 seconds
   Forward delay                    : 15 seconds
   Message age                      : 0
   Number of topology changes       : 58
   Time since last topology change  : 14127 seconds
   Local parameters
```

```
    Bridge ID                    : 28672.00:90:69:0b:3f:d0
    Extended system ID           : 0
```

show spanning-tree bridge vlan-id (VSTP)

```
user@host> show spanning-tree bridge vlan-id 10
STP bridge parameters
Routing instance name               : GLOBAL
Enabled protocol                    : RSTP

 STP bridge parameters for VLAN 10
   Root ID                          : 28672.00:90:69:0b:3f:d0
   Hello time                       : 2 seconds
   Maximum age                      : 20 seconds
   Forward delay                    : 15 seconds
   Message age                      : 0
   Number of topology changes       : 58
   Time since last topology change  : 14127 seconds
   Local parameters
     Bridge ID                      : 28672.00:90:69:0b:3f:d0
     Extended system ID             : 0
```

show spanning-tree bridge

Syntax

show spanning-tree bridge
<brief | detail>
<msti *msti-id*>
<vlan *vlan-id*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display the configured or calculated spanning-tree protocol (can be either STP, RSTP, or MSTP) parameters.

Options

none—(Optional) Display brief STP bridge information for all Multiple Spanning Tree Instances (MSTIs).

brief | detail—(Optional) Display the specified level of output.

msti *msti-id*—(Optional) Display STP bridge information for the specified MSTP instance ID or Common and Internal Spanning Tree (CIST). Specify **0** for CIST. Specify a value from **1** through **4094** for an MSTI.

vlan *vlan-id*—(Optional) Display STP bridge information for the specified VLAN. Specify a VLAN tag identifier from **1** through **4094**.

Required Privilege Level

view

Related Documentation

- **show spanning-tree interface on page 1407**
- Example: Configuring Network Regions for VLANs with MSTP on J-EX Series Switches on page 1297
- Understanding STP for J-EX Series Switches on page 1275
- Understanding RSTP for J-EX Series Switches on page 1276
- Understanding MSTP for J-EX Series Switches on page 1277
- Understanding VSTP for J-EX Series Switches on page 1281

List of Sample Output

show spanning-tree bridge on page 1400
show spanning-tree bridge brief on page 1400
show spanning-tree bridge detail on page 1401

Output Fields

Table 180 on page 1398 lists the output fields for the **show spanning-tree bridge** command. Output fields are listed in the approximate order in which they appear.

Table 180: show spanning-tree bridge Output Fields

Field Name	Field Description
Context ID	An internally generated identifier.

Table 180: show spanning-tree bridge Output Fields *(continued)*

Field Name	Field Description
Enabled protocol	Spanning-tree protocol type enabled.
Root ID	Bridge ID of the elected spanning tree root bridge. The bridge ID consists of a configurable bridge priority and the MAC address of the bridge.
Root cost	Calculated cost to reach the root bridge from the bridge where the command is entered.
Root port	Interface that is the current elected root port for this bridge.
CIST regional root	Bridge ID of the elected MSTP regional root bridge.
CIST internal root cost	Calculated cost to reach the regional root bridge from the bridge where the command is entered.
Hello time	Configured number of seconds between transmissions of configuration BPDUs.
Maximum age	Maximum age of received protocol BPDUs.
Forward delay	Configured time an STP bridge port remains in the listening and learning states before transitioning to the forwarding state.
Hop count	Configured maximum number of hops a BPDU can be forwarded in the MSTP region.
Message age	Number of seconds elapsed since the most recent BPDU was received.
Number of topology changes	Total number of STP topology changes detected since the switch last booted.
Time since last topology change	Number of seconds elapsed since the most recent topology change.
Topology change initiator	Interface name of the interface that received the topology change request.
Topology change last recvd. from	Bridge ID of the bridge that requested the last topology change.
Bridge ID (Local)	Locally configured bridge ID. The bridge ID consists of a configurable bridge priority and the MAC address of the bridge.
Extended system ID	Internally generated system identifier.
MSTI regional root	Bridge ID of the elected MSTP regional root bridge.
Internal instance ID	An internally generated identifier.

Table 180: show spanning-tree bridge Output Fields *(continued)*

Field Name	Field Description
Path Cost Method	Bridges supporting 802.1D (legacy) implement only 16-bit values for path cost. Newer versions of this standard support 32-bit values.

show spanning-tree bridge

```
user@switch> show spanning-tree bridge
 STP bridge parameters
 Context ID                          : 0
 Enabled protocol                    : MSTP

 STP bridge parameters for CIST
   Root ID                           : 32768.00:11:f2:56:df:40
   Root cost                         : 0
   Root port                         : ge-0/0/1.0
   CIST regional root                : 32768.00:11:f2:56:df:40
   CIST internal root cost           : 20000
   Hello time                        : 2 seconds
   Maximum age                       : 20 seconds
   Forward delay                     : 15 seconds
   Hop count                         : 19
   Message age                       : 0
   Number of topology changes        : 1
   Time since last topology change   : 108 seconds
   Topology change initiator         : ge-0/0/1.0
   Topology change last recvd. from  : 00:11:f2:56:df:4c
   Local parameters
     Bridge ID                       : 32768.00:11:f2:57:1c:00
     Extended system ID              : 0
     Internal instance ID            : 0

 STP bridge parameters for MSTI 10
   MSTI regional root                : 32778.00:11:f2:56:df:40
   Root cost                         : 20000
   Root port                         : ge-0/0/1.0
   Hello time                        : 2 seconds
   Maximum age                       : 20 seconds
   Forward delay                     : 15 seconds
   Hop count                         : 19
   Number of topology changes        : 1
   Time since last topology change   : 108 seconds
   Topology change initiator         : ge-0/0/1.0
   Topology change last recvd. from  : 00:11:f2:56:df:41
   Local parameters
     Bridge ID                       : 32778.00:11:f2:57:1c:00
     Extended system ID              : 0
     Internal instance ID            : 1
```

show spanning-tree bridge brief

```
user@switch> show spanning-tree bridge brief
STP bridge parameters
Context ID      : 0
Enabled protocol    : RSTP
  Root ID       : 32768.00:19:e2:50:95:a0
  Hello time        : 2 seconds
  Maximum age       : 20 seconds
  Forward delay     : 15 seconds
  Message age       : 0
```

```
  Number of topology changes    : 0
  Local parameters
    Bridge ID       : 32768.00:19:e2:50:95:a0
    Extended system ID      : 0
    Internal instance ID    : 0
```

show spanning-tree bridge detail

```
user@switch> show spanning-tree bridge detail

STP bridge parameters
Context ID      : 0
Enabled protocol    : RSTP
  Root ID       : 32768.00:19:e2:50:95:a0
  Hello time        : 2 seconds
  Maximum age       : 20 seconds
  Forward delay     : 15 seconds
  Message age       : 0
  Number of topology changes    : 0
  Local parameters
    Bridge ID       : 32768.00:19:e2:50:95:a0
    Extended system ID      : 0
    Internal instance ID    : 0
    Hello time      : 2 seconds
    Maximum age     : 20 seconds
    Forward delay       : 15 seconds
    Path cost method    : 32 bit
```

show spanning-tree interface

Syntax

show spanning-tree interface
<brief | detail>
<msti *msti-id*>
<routing-instance *routing-instance-name*>
<vlan-id *vlan-id*>

Syntax (J-EX Series Switch)

show spanning-tree interface
<brief | detail>
<msti *msti-id*>
<vlan-id *vlan-id*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display the configured or calculated interface-level STP parameters.

Options

none—Display brief STP interface information.

brief | detail—(Optional) Display the specified level of output.

msti *msti-id*—(Optional) Display STP interface information for the specified MST instance.

routing-instance *routing-instance-name*—(Optional) Display STP interface information for the specified routing instance.

vlan-id *vlan-id*—(Optional) Display STP interface information for the specified VLAN.

Required Privilege Level

view

List of Sample Output

show spanning-tree interface on page 1403
show spanning-tree interface detail on page 1404
show spanning-tree interface msti on page 1405
show spanning-tree interface vlan-id 101 on page 1406
show spanning-tree interface (VSTP) on page 1406
show spanning-tree interface vlan-id (VSTP) on page 1406

Output Fields

Table 181 on page 1402 lists the output fields for the **show spanning-tree Interface** command. Output fields are listed in the approximate order in which they appear.

Table 181: show spanning-tree Interface Output Fields

Field Name	Field Description
Interface name	Interface configured to participate in the STP, RSTP, VSTP, or MSTP instance.
Port ID	Logical interface identifier configured to participate in the MSTP or VSTP instance.
Designated port ID	Port ID of the designated port for the LAN segment to which this interface is attached.

Table 181: show spanning-tree Interface Output Fields *(continued)*

Field Name	Field Description
Designated bridge ID	Bridge ID of the designated bridge for the LAN segment to which this interface is attached.
Port Cost	Configured cost for the interface.
Port State	STP port state: forwarding (**FWD**), blocking (**BLK**), listening, learning, or disabled.
Port Role	MSTP, VSTP, or RSTP port role: designated (**DESG**), backup (**BKUP**), alternate (**ALT**), root, or Root Prevented (**Root-Prev**).
Link type	MSTP, VSTP, or RSTP link type. Shared or point-to-point (pt-pt) and edge or nonedge.
Alternate	Identifies the interface as an MSTP, VSTP, or RSTP alternate root port (**Yes**) or nonalternate root port (**No**).
Boundary Port	Identifies the interface as an MSTP regional boundary port (**Yes**) or nonboundary port (**No**).

show spanning-tree interface

```
user@host> show spanning-tree interface routing-instance vs1 detail
Spanning tree interface parameters for instance 0

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID       bridge ID          Cost
ae1             128:1        128:1  32768.0090690b47d1      1000  FWD    DESG
ge-2/1/2        128:2        128:2  32768.0090690b47d1     20000  FWD    DESG
ge-2/1/5        128:3        128:3  32768.0090690b47d1     29999  FWD    DESG
ge-2/2/1        128:4       128:26  32768.0013c39ec880     20000  FWD    ROOT
xe-9/2/0        128:5        128:5  32768.0090690b47d1      2000  FWD    DESG
xe-9/3/0        128:6        128:6  32768.0090690b47d1      2000  FWD    DESG

Spanning tree interface parameters for instance 1

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID       bridge ID          Cost
ae1             128:1        128:1  32769.0090690b47d1      1000  FWD    DESG
ge-2/1/2        128:2        128:2  32769.0090690b47d1     20000  FWD    DESG
ge-2/1/5        128:3        128:3  32769.0090690b47d1     29999  FWD    DESG
ge-2/2/1        128:4       128:26  32769.0013c39ec880     20000  FWD    ROOT
xe-9/2/0        128:5        128:5  32769.0090690b47d1      2000  FWD    DESG
xe-9/3/0        128:6        128:6  32769.0090690b47d1      2000  FWD    DESG

Spanning tree interface parameters for instance 2

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID       bridge ID          Cost
ae1             128:1        128:1  32770.0090690b47d1      1000  FWD    DESG
ge-2/1/2        128:2        128:2  32770.0090690b47d1     20000  FWD    DESG
ge-2/1/5        128:3        128:3  32770.0090690b47d1     29999  FWD    DESG
ge-2/2/1        128:4       128:26  32770.0013c39ec880     20000  FWD    ROOT
```

```
xe-9/2/0         128:5        128:5  32770.0090690b47d1      2000  FWD    DESG
xe-9/3/0         128:6        128:6  32770.0090690b47d1      2000  FWD    DESG
```

show spanning-tree interface detail

```
user@host> show spanning-tree interface routing-instance vs1 detail
Spanning tree interface parameters for instance 0

Interface name                : ae1
Port identifier               : 128.1
Designated port ID            : 128.1
Port cost                     : 1000
Port state                    : Forwarding
Designated bridge ID          : 32768.00:90:69:0b:47:d1
Port role                     : Designated
Link type                     : Pt-Pt/NONEDGE
Boundary port                  : No

Interface name                : ge-2/1/2
Port identifier               : 128.2
Designated port ID            : 128.2
Port cost                     : 20000
Port state                    : Forwarding
Designated bridge ID          : 32768.00:90:69:0b:47:d1
Port role                     : Designated
Link type                     : Pt-Pt/NONEDGE
Boundary port                  : No

Interface name                : ge-2/1/5
Port identifier               : 128.3
Designated port ID            : 128.3
Port cost                     : 29999
Port state                    : Forwarding
Designated bridge ID          : 32768.00:90:69:0b:47:d1
Port role                     : Designated
Link type                     : Pt-Pt/NONEDGE
Boundary port                  : No

Interface name                : ge-2/2/1
Port identifier               : 128.4
Designated port ID            : 128.26
Port cost                     : 20000
Port state                    : Forwarding
Designated bridge ID          : 32768.00:13:c3:9e:c8:80
Port role                     : Root
Link type                     : Pt-Pt/NONEDGE
Boundary port                  : No

Interface name                : xe-9/2/0
Port identifier               : 128.5
Designated port ID            : 128.5
Port cost                     : 2000
Port state                    : Forwarding
Designated bridge ID          : 32768.00:90:69:0b:47:d1
Port role                     : Designated
Link type                     : Pt-Pt/NONEDGE
Boundary port                  : No

Interface name                : xe-9/3/0
Port identifier               : 128.6
Designated port ID            : 128.6
Port cost                     : 2000
Port state                    : Forwarding
```

```
Designated bridge ID           : 32768.00:90:69:0b:47:d1
Port role                      : Designated
Link type                      : Pt-Pt/NONEDGE
Boundary port                   : No


Spanning tree interface parameters for instance 1

Interface name                 : ae1
Port identifier                : 128.1
Designated port ID             : 128.1
Port cost                      : 1000
Port state                     : Forwarding
Designated bridge ID           : 32768.00:90:69:0b:47:d1
Port role                      : Designated
Link type                      : Pt-Pt/NONEDGE
Boundary port                   : No

Interface name                 : ge-2/1/2
Port identifier                : 128.2
Designated port ID             : 128.2
Port cost                      : 20000
Port state                     : Forwarding
Designated bridge ID           : 32768.00:90:69:0b:47:d1
Port role                      : Designated
Link type                      : Pt-Pt/NONEDGE
Boundary port                   : No

Interface name                 : ge-2/1/5
Port identifier                : 128.3
Designated port ID             : 128.3
Port cost                      : 29999
Port state                     : Forwarding
Designated bridge ID           : 32768.00:90:69:0b:47:d1
Port role                      : Designated
Link type                      : Pt-Pt/NONEDGE
Boundary port                   : No

Interface name                 : ge-2/2/1
Port identifier                : 128.4
Designated port ID             : 128.26
Port cost                      : 20000
Port state                     : Forwarding
Designated bridge ID           : 32768.00:13:c3:9e:c8:80
Port role                      : Root
Link type                      : Pt-Pt/NONEDGE
Boundary port                   : No

...
```

show spanning-tree interface msti

```
user@host> show spanning-tree interface msti 1 routing-instance vs1 detail
 Spanning tree interface parameters for instance 1

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
xe-7/0/0         128:1        128:1  32769.0090690b4fd1      2000  FWD    DESG
ge-5/1/0         128:2        128:2  32769.0090690b4fd1     20000  FWD    DESG
ge-5/1/1         128:3        128:3  32769.0090690b4fd1     20000  FWD    DESG
ae1              128:4        128:1  32769.0090690b47d1     10000  BLK    ALT
```

```
ge-5/1/4         128:5        128:3  32769.0090690b47d1     20000  BLK    ALT
xe-7/2/0         128:6        128:6  32769.0090690b47d1      2000  FWD    ROOT
```

show spanning-tree interface vlan-id 101

```
user@host> show spanning-tree interface vlan-id 101 routing-instance vs1 detail
Spanning tree interface parameters for instance 0

Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID         bridge ID          Cost
ge-11/0/5        128:1        128:1  32768.0090690b7fd1     20000  FWD    DESG
ge-11/0/6        128:2        128:1  32768.0090690b7fd1     20000  BLK    BKUP
ge-11/1/0        128:3        128:2  32768.0090690b4fd1     20000  BLK    ALT
ge-11/1/1        128:4        128:3  32768.0090690b4fd1     20000  BLK    ALT
ge-11/1/4        128:5        128:1  32768.0090690b47d1     20000  BLK    ALT
xe-10/0/0        128:6        128:5  32768.0090690b4fd1      2000  BLK    ALT
xe-10/2/0        128:7        128:4  32768.0090690b47d1      2000  FWD    ROOT
```

show spanning-tree interface (VSTP)

```
user@host> show spanning-tree interface
Spanning tree interface parameters for instance 0

Interface    Port ID    Designated      Designated         Cost    State  Role
                         port ID         bridge ID
ge-1/0/1         128:1        128:1  28672.0090690b3fe0     20000  FWD    DESG
ge-1/0/2         128:2        128:2  28672.0090690b3fe0     20000  FWD    DESG

Spanning tree interface parameters for VLAN 10

Interface    Port ID    Designated      Designated         Cost    State  Role
                         port ID         bridge ID
ge-1/0/1         128:1        128:1  28672.0090690b3fe0     20000  FWD    DESG
ge-1/0/2         128:2        128:2  28672.0090690b3fe0     20000  FWD    DESG

Spanning tree interface parameters for VLAN 20

Interface    Port ID    Designated      Designated         Cost    State  Role
                         port ID         bridge ID
ge-1/0/1         128:1        128:1  28672.0090690b3fe0     20000  FWD    DESG
ge-1/0/2         128:2        128:2  28672.0090690b3fe0     20000  FWD    DESG
```

show spanning-tree interface vlan-id (VSTP)

```
user@host> show spanning-tree interface vlan-id 10
Spanning tree interface parameters for VLAN 10

Interface    Port ID    Designated      Designated         Cost    State  Role
                         port ID         bridge ID
ge-1/0/1         128:1        128:1  28672.0090690b3fe0     20000  FWD    DESG
ge-1/0/2         128:2        128:2  28672.0090690b3fe0     20000  FWD    DESG
```

show spanning-tree interface

Syntax

show spanning-tree interface
<brief | detail>
<interface-name *interface-name*>
<msti *msti-id*>
<vlan-id *vlan-id*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display the configured or calculated interface-level spanning-tree protocol (can be either STP, RSTP, or MSTP) parameters. In **brief** mode, will not display interfaces that are administratively disabled or do not have a physical link.

Options

none—(Optional) Display brief STP interface information.

brief | detail—(Optional) Display the specified level of output.

interface-name *interface-name*—(Optional) Name of an interface.

msti *msti-id*—(Optional) Display STP bridge information for the specified MSTP instance ID or Common and Internal Spanning Tree (CIST). Specify **0** for CIST. Specify a value from **1** through **4094** for an MSTI.

vlan-id *vlan-id*—(Optional) For MSTP interfaces, display interface information for the specified VLAN. Specify a value from **0** through **4094**.

Required Privilege Level

view

Related Documentation

- **show spanning-tree bridge on page 1398**
- Example: Configuring Network Regions for VLANs with MSTP on J-EX Series Switches on page 1297
- Understanding STP for J-EX Series Switches on page 1275
- Understanding RSTP for J-EX Series Switches on page 1276
- Understanding MSTP for J-EX Series Switches on page 1277
- Understanding VSTP for J-EX Series Switches on page 1281

List of Sample Output

show spanning-tree interface on page 1408
show spanning-tree interface brief on page 1409
show spanning-tree interface detail on page 1409
show spanning-tree interface ge-1/0/0 on page 1410

Output Fields

Table 182 on page 1408 lists the output fields for the **show spanning-tree Interface** command. Output fields are listed in the approximate order in which they appear.

Table 182: show spanning-tree interface Output Fields

Field Name	Field Description
Interface name	Interface configured to participate in the STP, RSTP, or MSTP instance.
Port ID	Logical interface identifier configured to participate in the MSTP instance.
Designated port ID	Port ID of the designated port for the LAN segment this interface is attached to.
Designated bridge ID	Bridge ID of the designated bridge for the LAN segment this interface is attached to.
Port Cost	Configured cost for the interface.
Port State	STP port state. Forwarding (FWD), blocking (BLK), listening, learning, or disabled.
Port Role	MSTP or RSTP port role. Designated (DESG), backup (BKUP), alternate (ALT), or root.
Link type	MSTP or RSTP link type. Shared or point-to-point (pt-pt) and edge or non edge.
Alternate	Identifies the interface as an MSTP or RSTP alternate root port (yes) or nonalternate root port (no).
Boundary Port	Identifies the interface as an MSTP regional boundary port (yes) or nonboundary port (no).

show spanning-tree interface

```
user@switch> show spanning-tree interface

Spanning tree interface parameters for instance 0

Interface    Port ID    Designated  Designated     Port    State  Role
                          port ID    bridge ID      Cost
ge-0/0/0.0  128:513  128:513   8192.0019e2500340    1000   FWD    DESG
ge-0/0/2.0  128:515  128:515   8192.0019e2500340    1000   BLK    DIS
ge-0/0/4.0  128:517  128:517   8192.0019e2500340    1000   FWD    DESG
ge-0/0/23.0 128:536  128:536   8192.0019e2500340    1000   FWD    DESG


Spanning tree interface parameters for instance 1

Interface    Port ID    Designated  Designated     Port    State  Role
                          port ID    bridge ID      Cost
ge-0/0/0.0  128:513  128:513   8193.0019e2500340    1000   FWD    DESG
ge-0/0/2.0  128:515  128:515   8193.0019e2500340    1000   BLK    DIS
ge-0/0/4.0  128:517  128:517   8193.0019e2500340    1000   FWD    DESG
ge-0/0/23.0 128:536  128:536   8193.0019e2500340    1000   FWD    DESG

Spanning tree interface parameters for instance 2

Interface    Port ID    Designated  Designated     Port    State  Role
                          port ID    bridge ID      Cost
ge-0/0/0.0  128:513  128:1   8194.001b549fd000  1000  FWD    ROOT
ge-0/0/2.0  128:515  128:515  32770.0019e2500340  4000  BLK    DIS
ge-0/0/4.0  128:517  128:1    16386.001b54013080  1000  BLK    ALT
```

```
ge-0/0/23.0 128:536  128:536  32770.0019e2500340  1000  FWD    DESG
```

show spanning-tree interface brief

```
user@switch> show spanning-tree interface brief
Spanning tree interface parameters for instance 0

Interface    Port ID    Designated  Designated     Port    State  Role
     port ID    bridge ID  Cost
ge-1/0/0.0  128:625  128:625  32768.0019e25095a0 20000  BLK    DIS
ge-1/0/1.0  128:626  128:626  32768.0019e25095a0 20000  BLK    DIS
ge-1/0/2.0  128:627  128:627  32768.0019e25095a0 20000  BLK    DIS
ge-1/0/10.0 128:635  128:635  32768.0019e25095a0 20000  BLK    DIS
ge-1/0/20.0 128:645  128:645  32768.0019e25095a0 20000  BLK    DIS
ge-1/0/30.0 128:655  128:655  32768.0019e25095a0 20000  BLK    DIS
```

show spanning-tree interface detail

```
user@switch> show spanning-tree interface detail
Spanning tree interface parameters for instance 0

Interface name     : ge-1/0/0.0
Port identifier    : 128.625
Designated port ID     : 128.625
Port cost      : 20000
Port state     : Blocking
Designated bridge ID   : 32768.00:19:e2:50:95:a0
Port role      : Disabled
Link type      : Pt-Pt/EDGE
Boundary port       : NA

Interface name     : ge-1/0/1.0
Port identifier    : 128.626
Designated port ID     : 128.626
Port cost      : 20000
Port state     : Blocking
Designated bridge ID   : 32768.00:19:e2:50:95:a0
Port role      : Disabled
Link type      : Pt-Pt/NONEDGE
Boundary port       : NA

Interface name     : ge-1/0/2.0
Port identifier    : 128.627
Designated port ID     : 128.627
Port cost      : 20000
Port state     : Blocking
Designated bridge ID   : 32768.00:19:e2:50:95:a0
Port role      : Disabled
Link type      : Pt-Pt/NONEDGE
Boundary port       : NA

Interface name     : ge-1/0/10.0
Port identifier    : 128.635
Designated port ID     : 128.635
Port cost      : 20000
Port state     : Blocking
Designated bridge ID   : 32768.00:19:e2:50:95:a0
Port role      : Disabled
Link type      : Pt-Pt/NONEDGE
Boundary port       : NA

Interface name     : ge-1/0/20.0
Port identifier    : 128.645
Designated port ID     : 128.645
```

```
Port cost      : 20000
Port state     : Blocking
Designated bridge ID   : 32768.00:19:e2:50:95:a0
Port role      : Disabled
Link type      : Pt-Pt/NONEDGE
Boundary port       : NA
[output truncated]
```

show spanning-tree interface ge-1/0/0

```
user@switch> show spanning-tree interface ge-1/0/0

Interface    Port ID    Designated  Designated     Port    State  Role
     port ID    bridge ID  Cost
ge-1/0/0.0 128:625  128:625  32768.0019e25095a0 20000  BLK    DIS
```

show spanning-tree mstp configuration

Syntax
show spanning-tree mstp configuration
<brief | detail>
<routing-instance *routing-instance-name*>

Syntax (J-EX Series Switch)
show spanning-tree mstp configuration
<brief | detail>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display the MSTP configuration.

Options
none—Display MSTP configuration information.

brief | detail—(Optional) Display the specified level of output.

routing-instance *routing-instance-name*—(Optional) Display MSTP configuration information for the specified routing instance.

Required Privilege Level
view

List of Sample Output
show spanning-tree mstp configuration on page 1411

Output Fields
Table 183 on page 1411 lists the output fields for the **show spanning-tree mstp configuration** command. Output fields are listed in the approximate order in which they appear.

Table 183: show spanning-tree mstp configuration Output Fields

Field Name	Field Description
Context id	Internally generated identifier.
Region name	MSTP region name carried in the MSTP BPDUs.
Revision	Revision number of the MSTP configuration.
Configuration digest	Numerical value derived from the VLAN-to-instance mapping table.
MSTI ID	MST instance identifier.
Member VLANs	VLAN identifiers associated with the MSTI.

show spanning-tree mstp configuration

```
user@host> show spanning-tree mstp configuration routing-instance vs1 detail
MSTP configuration information
Context identifier     : 1
Region name            : henry
Revision               : 3
Configuration digest   : 0x6da4b5c4fd587757eef35675365e1
```

```
MSTI     Member VLANs
   0 0-99,101-199,201-4094
   1 100
   2 200
```

show spanning-tree mstp configuration

Syntax show spanning-tree mstp configuration
<brief | detail>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display the MSTP configuration.

Options none—Display MSTP configuration information.

brief | detail—(Optional) Display the specified level of output.

Required Privilege Level view

List of Sample Output **show spanning-tree mstp configuration on page 1413**

Output Fields Table 184 on page 1413 lists the output fields for the **show spanning-tree mstp configuration** command. Output fields are listed in the approximate order in which they appear.

Table 184: show spanning-tree mstp configuration Output Fields

Field Name	Field Description
Context identifier	Internally generated identifier.
Region name	MSTP region name carried in the MSTP BPDUs.
Revision	Revision number of the MSTP configuration.
Configuration digest	Numerical value derived from the VLAN-to-instance mapping table.
MSTI	MSTI instance identifier.
Member VLANs	Identifiers for VLANs associated with the MSTI.

show spanning-tree mstp configuration

```
user@host> show spanning-tree mstp configuration
MSTP configuration information
Context identifier     : 0
Region name            : region1
Revision               : 0
Configuration digest   : 0xc92e7af9febb44d8df928b87f16b


MSTI     Member VLANs
   0 0-100,105-4094
   1 101-102
   2 103-104
```

show spanning-tree statistics

Syntax

show spanning-tree statistics
<brief | detail >
<interface *interface-name*>
<routing-instance *routing-instance-name*>

Syntax (J-EX Series Switch)

show spanning-tree statistics
<brief | detail >
<interface *interface-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display STP statistics.

Options

none—Display brief STP statistics.

brief | detail—(Optional) Display the specified level of output.

interface *interface-name*—(Optional) Display STP statistics for the specified interface.

routing-instance *routing-instance-name*—(Optional) Display STP statistics for the specified routing instance.

Required Privilege Level

view

List of Sample Output

show spanning-tree statistics routing-instance on page 1415
show spanning-tree statistics interface on page 1415

Output Fields

Table 185 on page 1414 lists the output fields for the **show spanning-tree statistics** command. Output fields are listed in the approximate order in which they appear.

Table 185: show spanning-tree statistics Output Fields

Field Name	Field Description
Message type	Type of message being counted.
BPDUs sent	Total number of BPDUs sent.
BPDUs received	Total number of BPDUs received.
BPDUs sent in last 5 secs	Number of BPDUs sent in the most recent 5-second period.
BPDUs received in last 5 secs	Number of BPDUs received in the most recent 5-second period.
Interface	Interface for which the statistics are being displayed.
Next BPDU transmission	Number of seconds until the next BPDU is scheduled to be sent.

show spanning-tree statistics routing-instance

```
user@host> show spanning-tree statistics routing-instance vs1 detail
Routing instance level STP statistics
Message type                   : bpdus
BPDUs sent                     : 121
BPDUs received                 : 537
BPDUs sent in last 5 secs      : 5
BPDUs received in last 5 secs  : 27
```

show spanning-tree statistics interface

```
user@host> show spanning-tree statistics interface ge-11/1/4 routing-instance vs1 detail
Interface   BPDUs sent   BPDUs received      Next BPDU
                                            transmission
ge-11/1/4            7           190                 0
```

show spanning-tree statistics

Syntax

show spanning-tree statistics
interface *interface-name*
vlan *vlan-id*
<brief | detail>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display STP statistics on an interface, or for a VLAN when VSTP is enabled.

Options

none—Display brief STP statistics.

brief | detail—(Optional) Display the specified level of output.

interface *interface-name*—(Optional) The name of the interface.

vlan *vlan-id*—(Optional) The name of a VLAN.

Required Privilege Level

view

Related Documentation

- **show spanning-tree bridge on page 1398**
- Example: Configuring Network Regions for VLANs with MSTP on J-EX Series Switches on page 1297
- Understanding STP for J-EX Series Switches on page 1275
- Understanding RSTP for J-EX Series Switches on page 1276
- Understanding MSTP for J-EX Series Switches on page 1277
- Understanding VSTP for J-EX Series Switches on page 1281

List of Sample Output

show spanning-tree statistics interface on page 1417

Output Fields

Table 186 on page 1416 lists the output fields for the **show spanning-tree statistics** command. Output fields are listed in the approximate order in which they appear.

Table 186: show spanning-tree statistics Output Fields

Field Name	Field Description
BPDUs sent	Total number of BPDUs sent.
BPDUs received	Total number of BPDUs received.
Interface	Interface for which the statistics are being displayed.
Next BPDU transmission	Number of seconds until the next BPDU is scheduled to be sent.

show spanning-tree statistics interface

```
user@switch> show spanning-tree statistics interface ge-0/0/4
Interface   BPDUs sent   BPDUs received  Next BPDU
                transmission
ge-0/0/4      7   190    0
```

PART 15

Layer 3 Protocols

CHAPTER 70

Layer 3 Protocols—Overview

Layer 3 Protocols Supported on J-EX Series Switches

J-EX Series switches support the Junos OS Layer 3 features and configuration statements listed in Table 187 on page 1421:

Table 187: Supported Junos OS Layer 3 Protocol Statements and Features

Protocol	Notes	For More Information
BGP	Fully supported.	See the *Junos OS Routing Protocols Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.
BFD	Fully supported.	See the *Junos OS Routing Protocols Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.
ICMP	Fully supported.	See the *Junos OS Routing Protocols Configuration Guide* at http://www.juniper.net/techpubs/software/junos/
IGMPv1, v2 and v3	Fully supported.	See the *Junos OS Multicast Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.
IS-IS	Supported, with the exceptions noted in "Layer 3 Protocols Not Supported on J-EX Series Switches" on page 14.	See the *Junos OS Routing Protocols Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.
MPLS	Supported, with the exceptions noted in "Layer 3 Protocols Not Supported on J-EX Series Switches" on page 14.	See the *Junos OS MPLS Applications Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.

Table 187: Supported Junos OS Layer 3 Protocol Statements and Features *(continued)*

Protocol	Notes	For More Information
OSPFv1, v2 and v3	Supported, with the exceptions noted in "Layer 3 Protocols Not Supported on J-EX Series Switches" on page 14.	See the *Junos OS Routing Protocols Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.
PIM	Supported, with the exception of IPv6.	See the *Junos OS Multicast Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.
RIP	Fully supported.	See the *Junos OS Routing Protocols Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.
RIPng	Fully supported.	See the *Junos OS Routing Protocols Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.
SNMP	Fully supported.	See the *Junos OS Network Management Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.
VRRP	Fully supported with exception of IPv6 support of VRRP on routed VLAN interfaces (RVIs).	See "Understanding VRRP on J-EX Series Switches" on page 1425. See also the *Junos OS High Availability Guide* at http://www.juniper.net/techpubs/software/junos/.

Related Documentation

- Layer 3 Protocols Not Supported on J-EX Series Switches on page 14
- J-EX Series Switch Software Features Overview on page 3

Layer 3 Protocols Not Supported on J-EX Series Switches

J-EX Series switches do not support the Junos OS Layer 3 protocols and features listed in Table 188 on page 1422:

Table 188: Junos OS Layer 3 Protocol Statements and Features That Are Not Supported

Feature	Configuration Statements Not Supported on J-EX Series Switches
DVMRP	• **dvmrp** and subordinate statements
Flow aggregation (cflowd)	• **cflow** and subordinate statements
GRE	• Not supported
IPsec	• **[edit services]** statements related to IPsec

Table 188: Junos OS Layer 3 Protocol Statements and Features That Are Not Supported *(continued)*

Feature	Configuration Statements Not Supported on J-EX Series Switches
IS-IS: • ES-IS • IPv6 in multicast routing protocols	• **clns-routing** statement • **ipv6-multicast** statement • **lsp-interval** statement • **label-switched-path** statement • **lsp-lifetime** statement • **te-metric** statement
Logical routers	• **logical-routers** and subordinate statements
MLD	• **mld** and all subordinate statements
MPLS: • Fast Reroute (FRR) • Label Distribution Protocol (LDP) • Layer 3 VPNs • Multiprotocol BGP (MP-BGP) for VPN-IPv4 family • Pseudowire emulation (PWE3) • Routing policy statements related to Layer 3 VPNs and MPLS • Virtual Private LAN Service (VPLS)	• **ldp** and all subordinate statements
Network Address Translation (NAT)	• **nat** and subordinate statements • Policy statements related to NAT
OSPF	• **demand-circuit** statement • **label-switched-path** and subordinate statements • **neighbor** statement within an OSPF area • **peer-interface** and subordinate statements within an OSPF area • **sham-link** statement • **te-metric** statement
PIM: • IPv6	• **inet6** family
Routing instances: • Routing instance forwarding	• **l2vpn** and subordinate statements • **ldp** and subordinate statements • **vpls** and subordinate statements
SAP and SDP	• **sap** and all subordinate statements

Table 188: Junos OS Layer 3 Protocol Statements and Features That Are Not Supported *(continued)*

Feature	Configuration Statements Not Supported on J-EX Series Switches
General routing options in the **routing-options** hierarchy: • MPLS and label-switched-paths	• **auto-export** and subordinate statements • **dynamic-tunnels** and subordinate statements • **lsp-next-hop** and subordinate statements • **multicast** and subordinate statements • **p2mp-lsp-next-hop** and subordinate statements • **route-distinguisher-id** statement

Traffic sampling and fowarding in the **forwarding-options** hierarchy	• **accounting** and subordinate statements
	• **family mpls** and **family multiservice** under **hash-key** hierarchy
	• Under **monitoring** *group-name* **family inet output** hierarchy:
	  ◦ **cflowd** statement
	  ◦ **export-format-cflowd-version-5** statement
	  ◦ **flow-active-timeout** statement
	  ◦ **flow-export-destination** statement
	  ◦ **flow-inactive-timeout** statement
	  ◦ **interface** statement
	• **port-mirroring** statement (On J-EX Series switches, port mirroring is implemented using the **analyzer** statement.)
	• **sampling** and subordinate statements

Related Documentation

- Layer 3 Protocols Supported on J-EX Series Switches on page 13
- J-EX Series Switch Software Features Overview on page 3

Understanding Distributed Periodic Packet Management on J-EX Series Switches

Periodic packet management (PPM) is responsible for processing a variety of time-sensitive periodic tasks for particular processes so that other processes on the J-EX Series Switch can more optimally direct their resources. PPM is responsible for the periodic transmission of packets on behalf of its various client processes, which include the process that controls Link Aggregation Control Protocol (LACP), and also for receiving packets on behalf of these client processes. PPM also gathers some statistics and sends process-specific packets. PPM cannot be disabled and is always running on any operational switch.

The responsibility for PPM processing on the switch is distributed between the Routing Engine and either the access interfaces (on J-EX4200 switches) or the line cards (on J-EX8200 switches) for all protocols that use PPM by default. This distributed model provides a faster response time for protocols that use PPM than the response time provided by the nondistributed model.

If distributed PPM is disabled, the PPM process runs on the Routing Engine only.

Distributed PPM can be disabled for all protocols that use PPM or for a single protocol that uses PPM. There is no way to disable PPM entirely.

BEST PRACTICE: We recommend that, generally, you disable distributed PPM only if Dell Support advises you to do so (see “Requesting Technical Support” on page lxxi). You should disable distributed PPM only if you have a compelling reason to disable it.

Related Documentation

- Configuring Distributed Periodic Packet Management on a J-EX Series Switch (CLI Procedure) on page 1451

Understanding VRRP on J-EX Series Switches

J-EX Series Switches support the Virtual Router Redundancy Protocol (VRRP) and VRRP for IPv6. This topic covers:

Overview of VRRP on J-EX Series Switches

You can configure the Virtual Router Redundancy Protocol (VRRP) or VRRP for IPv6 on Gigabit Ethernet interfaces, 10-Gigabit Ethernet interfaces, and logical interfaces on J-EX Series switches. When VRRP is configured, the switches act as virtual routing platforms. VRRP enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts. The VRRP routing platforms share the IP address corresponding to the default route configured on the hosts. At any time, one of the VRRP routing platforms is the master (active) and the others are backups. If the master routing platform fails, one of the backup routing platforms becomes the new master, providing a virtual default routing platform and enabling traffic on the LAN to be routed without relying on a single routing platform. Using VRRP, a backup J-EX Series switch can take over a failed default switch within a few seconds. This is done with minimum loss of VRRP traffic and without any interaction with the hosts.

VRRP for IPv6 provides a much faster switchover to an alternate default routing platform than IPv6 Neighbor Discovery (ND) procedures. VRRP for IPv6 does not support the **authentication-type** or **authentication-key** statements.

NOTE: Do not confuse the VRRP master and backup routing platforms with the master and backup member switches of a Virtual Chassis configuration. The master and backup members of a Virtual Chassis configuration compose a single host. In a VRRP topology, one host operates as the master routing platform and another operates as the backup routing platform, as shown in Figure 40 on page 1427.

Switches running VRRP dynamically elect master and backup routing platforms. You can also force assignment of master and backup routing platforms using priorities from **1** through **255**, with **255** being the highest priority. In VRRP operation, the default master

routing platform sends advertisements to backup routing platforms at regular intervals. The default interval is 1 second. If the backup routing platforms do not receive an advertisement for a set period, the backup routing platform with the highest priority takes over as master and begins forwarding packets.

NOTE: Priority 255 cannot be set for routed VLAN interfaces (RVIs).

VRRP is defined in RFC 3768, *Virtual Router Redundancy Protocol*.

Examples of VRRP Topologies

Figure 39 on page 1426 illustrates a basic VRRP topology with J-EX Series switches. In this example, Switches A, B, and C are running VRRP and together they make up a virtual routing platform. The IP address of this virtual routing platform is **10.10.0.1** (the same address as the physical interface of Switch A).

Figure 39: Basic VRRP on J-EX Series Switches

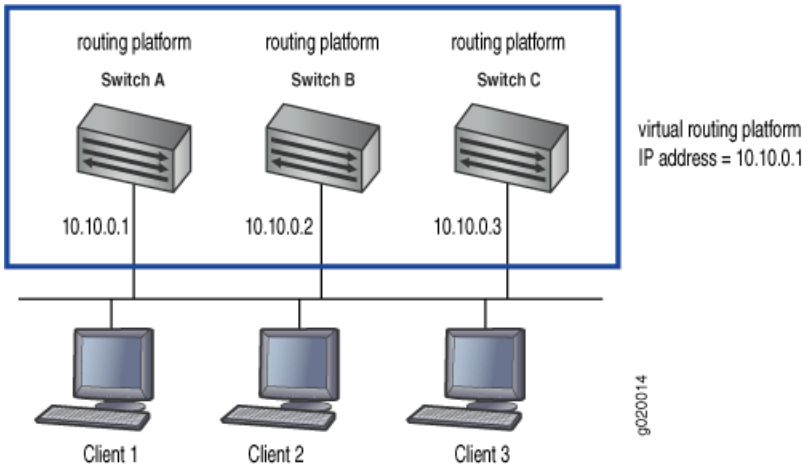

Figure 40 on page 1427 illustrates a basic VRRP topology using Virtual Chassis configurations. Switch A, Switch B, and Switch C are each composed of multiple interconnected J-EX4200 Ethernet Switches. Each Virtual Chassis configuration operates as a single switch, which is running VRRP, and together they make up a virtual routing platform. The IP address of this virtual routing platform is **10.10.0.1** (the same address as the physical interface of Switch A).

Figure 40: VRRP on Virtual Chassis Switches

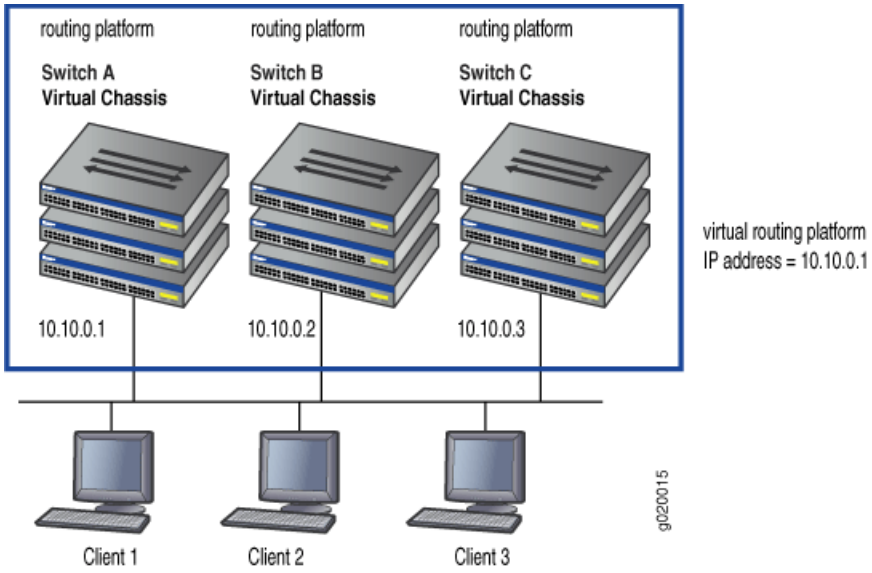

Because the virtual routing platform uses the IP address of the physical interface of Switch A, Switch A is the master VRRP routing platform, while Switch B and Switch C function as backup VRRP routing platforms. Clients 1 through 3 are configured with the default gateway IP address of **10.10.0.1** as the master router, Switch A, forwards packets sent to its IP address. If the master routing platform fails, the switch configured with the higher priority becomes the master virtual routing platform and provides uninterrupted service for the LAN hosts. When Switch A recovers, it becomes the master virtual routing platform again.

Related Documentation

- For more information on VRRP or VRRP for IPv6, see the *Junos OS High Availability Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.
- High Availability Features for J-EX Series Switches Overview on page 18
- Configuring VRRP for IPv6 (CLI Procedure) on page 1452

Understanding IPsec Authentication for OSPF Packets on J-EX Series Switches

IP Security (IPsec) provides a secure way to authenticate senders and encrypt IP version 4 (IPv4) traffic between network devices. IPsec offers network administrators for J-EX Series Ethernet Switches and their users the benefits of data confidentiality, data integrity, sender authentication, and anti-replay services.

IPsec is a framework for ensuring secure private communication over IP networks and is based on standards developed by the International Engineering Task Force (IETF). IPsec provides security services at the network layer of the Open Systems Interconnection (OSI) model by enabling a system to select required security protocols, determine the algorithms to use for the security services, and implement any cryptographic keys required to provide the requested services. You can use IPsec to protect one or more paths between a pair of hosts, between a pair of security gateways (such as switches), or between a security gateway and a host.

OSPF version 3 (OSPFv3), unlike OSPF version 2 (OSPFv2), does not have a built-in authentication method and relies on IPsec to provide this functionality. You can secure specific OSPFv3 interfaces and protect OSPFv3 virtual links.

Authentication Algorithms

Authentication is the process of verifying the identity of the sender. Authentication algorithms use a shared key to verify the authenticity of the IPsec devices. The Junos operating system (Junos OS) uses the following authentication algorithms:

- Message Digest 5 (MD5) uses a one-way hash function to convert a message of arbitrary length to a fixed-length message digest of 128 bits. Because of the conversion process, it is mathematically infeasible to calculate the original message by computing it backwards from the resulting message digest. Likewise, a change to a single character in the message will cause it to generate a very different message digest number.

  To verify that the message has not been tampered with, Junos OS compares the calculated message digest against a message digest that is decrypted with a shared key. Junos OS uses the MD5 hashed message authentication code (HMAC) variant that provides an additional level of hashing. MD5 can be used with an authentication header (AH) and Encapsulating Security Payload (ESP).

- Secure Hash Algorithm 1 (SHA-1) uses a stronger algorithm than MD5. SHA-1 takes a message of less than 264 bits in length and produces a 160-bit message digest. The large message digest ensures that the data has not been changed and that it originates from the correct source. Junos OS uses the SHA-1 HMAC variant that provides an additional level of hashing. SHA-1 can be used with AH, ESP, and Internet Key Exchange (IKE).

- SHA-256, SHA-384, and SHA-512 (sometimes grouped under the name SHA-2) are variants of SHA-1 and use longer message digests. Junos OS supports the SHA-256 version of SHA-2, which can process all versions of Advanced Encryption Standard (AES), Data Encryption Standard (DES), and Triple DES (3DES) encryption.

Encryption Algorithms

Encryption encodes data into a secure format so that it cannot be deciphered by unauthorized users. As with authentication algorithms, a shared key is used with encryption algorithms to verify the authenticity of IPsec devices. Junos OS uses the following encryption algorithms:

- Data Encryption Standard cipher-block chaining (DES-CBC) is a symmetric secret-key block algorithm. DES uses a key size of 64 bits, where 8 bits are used for error detection and the remaining 56 bits provide encryption. DES performs a series of simple logical operations on the shared key, including permutations and substitutions. CBC takes the first block of 64 bits of output from DES, combines this block with the second block, feeds this back into the DES algorithm, and repeats this process for all subsequent blocks.
- Triple DES-CBC (3DES-CBC) is an encryption algorithm that is similar to DES-CBC but provides a much stronger encryption result because it uses three keys for 168-bit (3 x 56-bit) encryption. 3DES works by using the first key to encrypt the blocks, the second key to decrypt the blocks, and the third key to reencrypt the blocks.

IPsec Protocols

IPsec protocols determine the type of authentication and encryption applied to packets that are secured by the switch. Junos OS supports the following IPsec protocols:

- AH—Defined in *RFC 2402*, AH provides connectionless integrity and data origin authentication for IPv4. It also provides protection against replays. AH authenticates as much of the IP header as possible, as well as the upper-level protocol data. However, some IP header fields might change in transit. Because the value of these fields might not be predictable by the sender, they cannot be protected by AH. In an IP header, AH can be identified with a value of 51 in the Protocol field of an IPv4 packet.
- ESP—Defined in *RFC 2406*, ESP can provide encryption and limited traffic flow confidentiality or connectionless integrity, data origin authentication, and an anti-replay service. In an IP header, ESP can be identified with a value of 50 in the Protocol field of an IPv4 packet.

Security Associations

An IPsec consideration is the type of security association (SA) that you wish to implement. An SA is a set of IPsec specifications that are negotiated between devices that are establishing an IPsec relationship. These specifications include preferences for the type of authentication, encryption, and IPsec protocol to be used when establishing the IPsec connection. An SA can be either unidirectional or bidirectional, depending on the choices made by the network administrator. An SA is uniquely identified by a Security Parameter

Index (SPI), an IPv4 or IPv6 destination address, and a security protocol (AH or ESP) identifier.

IPsec Modes

Junos OS supports the following IPsec modes:

- Tunnel mode is supported for both AH and ESP in Junos OS. In tunnel mode, the SA and associated protocols are applied to tunneled IPv4 or IPv6 packets. For a tunnel mode SA, an outer IP header specifies the IPsec processing destination and an inner IP header specifies the ultimate destination for the packet. The security protocol header appears after the outer IP header and before the inner IP header. In addition, there are slight differences for tunnel mode when you implement it with AH and ESP:
  - For AH, portions of the outer IP header are protected, as well as the entire tunneled IP packet.
  - For ESP, only the tunneled packet is protected, not the outer header.

  When one side of an SA is a security gateway (such as a switch), the SA must use tunnel mode. However, when traffic (for example, SNMP commands or BGP sessions) is destined for a switch, the system acts as a host. Transport mode is allowed in this case because the system does not act as a security gateway and does not send or receive transit traffic.

  

  NOTE: Tunnel mode is not supported for OSPF v3 control packet authentication.

- Transport mode provides an SA between two hosts. In transport mode, the protocols provide protection primarily for upper-layer protocols. A transport mode security protocol header appears immediately after the IP header and any options and before any higher-layer protocols (for example, TCP or UDP). There are slight differences for transport mode when you implement it with AH and ESP:
  - For AH, selected portions of the IP header are protected, as well as selected portions of the extension headers and selected options within the IPv4 header.
  - For ESP, only the higher-layer protocols are protected, not the IP header or any extension headers preceding the ESP header.

Related Documentation

- Using IP Security to Secure OSPFv3 Networks on page 1453
- Configuring an OSPF Network (J-Web Procedure) on page 1435

CHAPTER 71

Configuring Layer 3 Protocols

Configuring BGP Sessions (J-Web Procedure)

You can use the J-Web interface to create BGP peering sessions on a routing device.

NOTE: To configure BGP sessions, you must have a license for BGP installed on the J-EX Series switch.

To configure a BGP peering session:

1. Select **Configure** > **Routing** > **BGP**.

NOTE: After you make changes to the configuration in this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options** > **Commit**. See "Using the Commit Options to Commit Configuration Changes (J-Web Procedure)" on page 334 for details about all commit options.

2. Click one:

   - **Add**—Adds a BGP group. Enter information into the configuration page as described in Table 189 on page 1432.
   - **Edit**—Modifies an existing BGP group. Enter information into the configuration page as described in Table 189 on page 1432.
   - **Delete**—Deletes an existing BGP group.
   - **Disable**—Disables BGP configuration.

3. To modify BGP global settings, click **Edit** in the Global Information section. Enter information as described in Table 190 on page 1434.

Table 189: BGP Routing Configuration Summary

Field	Function	Your Action
General tab		
Group Type	Specifies whether the group is an internal BGP (IBGP) group or an external BGP (EBGP) group.	Select the option: **Internal** or **External**.
Group Name	Specifies the name for the group.	Type a new name or select and edit the name.
ASN	Sets the unique numeric identifier of the AS in which the routing device is configured.	Type the routing device's 32-bit AS number, in dotted decimal notation. If you enter an integer, the value is converted to a 32-bit equivalent. For example, if you enter **3**, the value assigned to the AS is **0.0.0.3**.
Preference	Specifies the degree of preference for an external route. The route with the highest local preference value is preferred.	Type or select and edit the value.
Cluster Id	Specifies the cluster identifier to be used by the route reflector cluster in an internal BGP group.	Type or select and edit the IPv6 or IPv4 address to be used as the identifier.
Description	Specifies the text description of the global, group, or neighbor configuration.	Type or select and edit the description.
Damping	Specifies whether route flap damping is enabled or not.	To enable route flap damping, select the check box. To disable route flap damping do not select the check box.
Advertise Inactive Routes	Specifies whether BGP advertises the best route even if the routing table did not select it to be an active route.	To enable advertising inactive routes, select the check box. To disable advertising inactive routes, do not select the check box.

Table 189: BGP Routing Configuration Summary *(continued)*

Field	Function	Your Action
Advertise Peer AS Routes	Specifies whether to disable the default behavior of suppressing AS routes.	To enable advertising peer AS routes, select the check box. To disable advertising peer AS routes, do not select the check box.
Neighbors tab		
Dynamic Neighbors	Configures a neighbor (peer).	Type the IPv4 address of the peer.
Static Neighbors	Configures the system's peers statically.	To configure a static neighbor: 1. Specify the IP address. 2. Specify the address of the local end of a BGP session. 3. Specify the degree of preference for an external route. 4. Enter a description. 5. Specify the hold-time value to use when negotiating a connection with the peer. 6. Specify how long a route must be present in the routing table before it is exported to BGP. Use this time delay to help bundle routing updates. 7. Select **Passive** if you do not want to send active open messages to the peer. 8. Select the option to compare the AS path of an incoming advertised route with the AS number of the BGP peer under the group and replace all occurrences of the peer AS number in the AS path with its own AS number before advertising the route to the peer. 9. Specify an import policy and export policy. 10. Click **OK**.
Policies tab		
Import Policy	Specifies one or more routing policies to routes being imported into the routing table from BGP.	Click **Add** to add an import policy. Select the policy and click **OK**. Click **Move up** or **Move down** to move the selected policy up or down the list of policies. Select the policy and click **Remove**.
Export Policy	Specifies one or more policies to routes being exported from the routing table into BGP.	Click **Add** to add an export policy. Select the policy and click **OK**. Click **Move up** or **Move down** to move the selected policy up or down the list of policies. Select the policy and click **Remove**.

Table 190: BGP Global Settings

Field	Function	Your Action
General tab		
Router ASN	Specifies the routing device's AS number.	Type or select and edit the value.
Router Identifier	Specify the routing device's IP address.	Type or select and edit the IP address.
BGP Status	Enables or disables BGP.	• To enable BGP, select **Enabled**. • To disable BGP, select **Disabled**.
Description	Describes of the global, group, or neighbor configuration.	Type or select and edit the description.
Confederation Number	Specifies the routing device's confederation AS number.	Type or select and edit the value.
Confederation Members	Specifies the AS numbers for the confederation members.	To add a member AS number, click **Add** and enter the number in the **Member ASN** box. Click **OK**. To modify a confederation member's AS number, select the member click **Edit** and, enter the number and click **OK**. To delete a confederation member, select the member and click **Remove**.
Advance Options	You can configure the following: • **Keep routes**—Specifies whether routes learned from a BGP peer must be retained in the routing table even if they contain an AS number that was exported from the local AS. • **TCP MSS**—Configures the maximum segment size (MSS) for the TCP connection for BGP neighbors. • **MTU Discovery**—Select to configure MTU discovery. • **Remove Private ASN**—Select to have the local system strip private AS numbers from the AS path when advertising AS paths to remote systems. • **Graceful Restart**—Specifies the time period when the restart is expected to be complete. Specify the maximum time that stale routes are kept during restart. • **Multihop**—Configures the maximum time-to-live (TTL) value for the TTL in the IP header of BGP packets. • **Authentication Type**—Select the authentication algorithm: None, MD5, SHA1, AES.	Select **All** or **None** to configure Keep Routes. Enter a value in the **TCP MSS** box. Click to enable **MTU Discovery**. Click to enable **Remove Private ASN**. Enter the time period for a graceful restart and the maximum time that stale routes must be kept. To configure Multihop, select **Nexthop Change** to allow unconnected third-party next hops. Enter a TTL value. Select the authentication algorithm. If you select None, specify an authentication key (password).
Policies tab		

Table 190: BGP Global Settings *(continued)*

Field	Function	Your Action
Import Policy	Specifies one or more routing policies to routes being imported into the routing table from BGP.	Click **Add** to add an import policy. Click **Move up** or **Move down** to move the selected policy up or down the list of policies. Click **Remove** to remove an import policy.
Export Policy	Specifies one or more policies to routes being exported from the routing table into BGP.	Click **Add** to add an export policy. Click **Move up** or **Move down** to move the selected policy up or down the list of policies. Click **Remove** to remove an export policy.
Trace Options tab		
File Name	Specifies the name of the file to receive the output of the tracing operation.	Type or select and edit the name.
Number of Files	Specifies the maximum number of trace files.	Type or select and edit the value.
File Size	Specifies the maximum size for each trace file.	Type or select and edit the value.
World Readable	Specifies whether the trace file can be read by any user or not.	Select **True** to allow any user to read the file. Select **False** to disallow all users being able to read the file.
Flags	Specifies the tracing operation to perform.	Select a value from the list.

Related Documentation

- Monitoring BGP Routing Information on page 1455
- Layer 3 Protocols Supported on J-EX Series Switches on page 13

Configuring an OSPF Network (J-Web Procedure)

You can use the J-Web interface to create multiarea OSPF networks on a J-EX Series switch.

To configure a multiarea OSPF network:

1. Select **Configure** > **Routing** > **OSPF**.

NOTE: After you make changes to the configuration in this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options** > **Commit**. See "Using the Commit Options to Commit Configuration Changes (J-Web Procedure)" on page 334 for details about all commit options.

2. Click one:
   - **Add**—Adds an OSPF area. Enter information into the configuration page as described in Table 191 on page 1436.
   - **Edit**—Modifies an existing OSPF area. Enter information into the configuration page as described in Table 191 on page 1436.
   - **Delete**—Deletes an existing OSPF area.
3. To modify OSPF global settings, click **Edit**. Enter information as described in Table 192 on page 1438.
4. To disable OSPF, click **Disable**.

Table 191: OSPF Routing Configuration Summary

Field	Function	Your Action
General tab		
Area Id	Uniquely identifies the area within its AS.	Type a 32-bit numeric identifier for the area. Type an integer or select and edit the value. If you enter an integer, the value is converted to a 32-bit equivalent. For example, if you enter 3, the value assigned to the area is **0.0.0.3**.
Area Ranges	Specifies a range of IP addresses for an area when sending summary link advertisements (within an area).	To add a range: 1. Click **Add**. 2. Type the area range. 3. Specify the subnet mask. 4. To override the metric for the IP address range, type a specific metric value. 5. If you do not want to display the routes that are contained within a summary, select **Restrict advertisements of this area range**. 6. If you want a summary of a route to be advertised only when an exact match is made with the configured summary range, select **Enforce exact match for advertisement of this area range**. 7. Click **OK**. To modify an existing area range, select the area range, click **Edit**, and edit the value. Click **OK**. To delete an area range, select the area range and click **Delete**.

Table 191: OSPF Routing Configuration Summary *(continued)*

Field	Function	Your Action
Area Type	Designates the type of OSPF area. • **regular**—A regular OSPF area, including the backbone area • **stub**—A stub area • **nssa**—A not-so-stubby area (NSSA)	Select the type of OSPF area you are creating from the list. If you select **stub**: 1. Enter the default metric. 2. To flood summary LSAs into the stub area, select the check box. If you select **nssa**: 1. Specify the metric type. 2. Enter the default metric. 3. To flood summary LSAs into the nssa area, select the check box. 4. To flood Type-7 LSAs into the nssa area, select the check box.
Interfaces tab		
Interfaces	Specifies the interfaces to be associated with the OSPF configuration	To associate an interface with the configuration, select the interface from the list, select **Associate** and click **OK**. To edit an interface's configuration: 1. Select the interface from the list and click **Edit**. 2. Specify the cost of an OSPF interface. 3. Specify the traffic engineering metric. 4. Specify how often the routing device sends hello packets from the interface. 5. Specify how long the routing device waits to receive a link-state acknowledgment packet before retransmitting link-state advertisements to an interface's neighbors. 6. To enable OSPF on the interface, select the check box. 7. To inform other protocols about neighbor down events, select the check box. 8. To treat the interface as a secondary interface, select the check box. 9. To only advertise OSPF, select the check box. 10. Click **OK**.
Policies tab		
Import Policy	Specifies one or more policies to control which routes learned from an area are used to generate summary link-state advertisements (LSAs) into other areas.	Click **Add** to add an import policy. Click **Move up** or **Move down** to move the selected policy up or down the list of policies. Click **Remove** to remove an import policy.

Table 191: OSPF Routing Configuration Summary *(continued)*

Field	Function	Your Action
Export Policy	Specifies one or more policies to control which summary LSAs are flooded into an area.	Click **Add** to add an export policy. Click **Move up** or **Move down** to move the selected policy up or down the list of policies. Click **Remove** to remove an export policy.

Table 192: Edit OSPF Global Settings

Field	Function	Your Action
General tab		
Router Id	Specifies the ID for the routing device.	Type or select and edit the value.
RIB Group	Installs the routes learned from OSPF routing instances into routing tables in the OSPF routing table group.	Select a value.
Internal Route Preference	Specifies the route preference for internal groups.	Type or select and edit the value.
External Route Preference	Specifies the route preference for external groups.	Type or select and edit the value.
Graceful Restart	Configures graceful restart for OSPF.	To configure graceful restart: 1. Specify the estimated time to send out purged grace LSAs over all the interfaces. 2. Specified the estimated time to reacquire a full OSPF neighbor from each area. 3. To disable **No Strict LSA Checking**, select the check box. 4. To disable graceful restart helper capability, select the check box. Helper mode is enabled by default. 5. Click **OK**.
SPF Options	Configure options for running the shortest-path-first (SPF) algorithm. You can configure a delay for when to run the SPF algorithm after a network topology change is detected, the maximum number of times the SPF algorithm can run in succession, and a hold-down interval after the SPF algorithm runs the maximum number of times.	To configure SPF: 1. Specify the time interval between the detection of a topology change and when the SPF algorithm runs. 2. Specify the time interval to hold down, or wait before a subsequent SPF algorithm runs after the SPF algorithm has run the configured maximum number of times in succession. 3. Specify the maximum number of times the SPF algorithm can run in succession. After the maximum is reached, the hold-down interval begins.

Table 192: Edit OSPF Global Settings *(continued)*

Field	Function	Your Action
Policies tab		
Import Policy	Specifies one or more policies to control which routes learned from an area are used to generate summary link-state advertisements (LSAs) into other areas.	Click **Add** to add an import policy. Click **Move up** or **Move down** to move the selected policy up or down the list of policies. Click **Remove** to remove an import policy.
Export Policy	Specifies one or more policies to control which summary LSAs are flooded into an area.	Click **Add** to add an export policy. Click **Move up** or **Move down** to move the selected policy up or down the list of policies. Click **Remove** to remove an export policy.
Trace Options tab		
File Name	Specifies the name of the file to receive the output of the tracing operation.	Type or select and edit the name.
Number of Files	Specifies the maximum number of trace files.	Type or select and edit the name.
File Size	Specifies the maximum size for each trace file.	Type or select and edit the name.
World Readable	Specifies whether the trace file can be read by any user or not.	Select **True** to allow any user to read the file. Select **False** to disallow all users being able to read the file.
Flags	Specifies the tracing operation to perform.	Select a value from the list.

Related Documentation

- Monitoring OSPF Routing Information on page 1457
- Layer 3 Protocols Supported on J-EX Series Switches on page 13

Configuring a RIP Network (J-Web Procedure)

You can use the J-Web interface to create RIP networks.

To configure a RIP network:

1. Select **Configure** > **Routing** > **RIP**.

NOTE: After you make changes to the configuration in this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options** > **Commit**. See "Using the Commit Options to Commit Configuration Changes (J-Web Procedure)" on page 334 for details about all commit options.

2. Click one:

   - **Add**—Configures a RIP instance. Enter information into the RIP Configuration page as described in Table 193 on page 1440.
   - **Edit**—Modifies an existing RIP instance. Enter information into the configuration page for RIP as described in Table 193 on page 1440.
   - **Delete**—Deletes an existing RIP instance.

4. To modify RIP global settings, click **Edit**. Enter information in the configuration as described in Table 194 on page 1441.

Table 193: RIP Routing Configuration Summary

Field	Function	Your Action
General tab		
Routing instance name	Specifies a name for the routing instance.	Type or select and edit the name.
Preference	Specifies the preference of external routes learned by RIP as compared to those learned from other routing protocols.	Type or select and edit the value.
Metric Out	Specifies the metric value to add to routes transmitted to the neighbor.	Type or select and edit the value.
Update interval	Specifies an update time interval to periodically send out routes learned by RIP to neighbors.	Type or select and edit the value.
Route timeout	Specifies the route timeout interval for RIP.	Type or select and edit the value.
Policies tab		
Import Policy	Applies one or more policies to routes being imported into the local routing device from the neighbors.	Click **Add** to add an import policy. Click **Move up** or **Move down** to move the selected policy up or down the list of policies. Click **Remove** to remove an import policy.

Table 193: RIP Routing Configuration Summary *(continued)*

Field	Function	Your Action
Export Policy	Applies a policy to routes being exported to the neighbors.	Click **Add** to add an export policy. Click **Move up** or **Move down** to move the selected policy up or down the list of policies. Click **Remove** to remove an export policy.
Neighbors tab		
RIP-Enabled Interfaces	Selects the interfaces to be associated with the RIP instance.	To enable RIP on an interface, click the check box next to the interface name. Click Edit if you want to modify an interface's settings.

Table 194: Edit RIP Global Settings

Field	Function	Your Action
General tab		
Send	Specifies RIP send options.	Select a value.
Receive	Configure RIP receive options.	Select a value.
Route timeout (sec)	Specifies the route timeout interval for RIP.	Type a value.
Update interval (sec)	Specifies the update time interval to periodically send out routes learned by RIP to neighbors.	Type or select and edit the value.
Hold timeout (sec)	Specifies the time period the expired route is retained in the routing table before being removed.	Type or select and edit the value.
Metric in	Specifies the metric to add to incoming routes when advertising into RIP routes that were learned from other protocols.	Type or select and edit the value.
RIB Group	Specifies a routing table group to install RIP routes into multiple routing tables.	Select and edit the name of the routing table group.
Message size	Specifies the number of route entries to be included in every RIP update message.	Type or select and edit the value.

Table 194: Edit RIP Global Settings *(continued)*

Field	Function	Your Action
Check Zero	Specifies whether the reserved fields in a RIP packet are zero. Options are: • **check-zero**—Discard version 1 packets that have nonzero values in the reserved fields and version 2 packets that have nonzero values in the fields that must be zero. This default behavior implements the RIP version 1 and version 2 specifications. • **no-check-zero**—Receive RIP version 1 packets with nonzero values in the reserved fields or RIP version 2 packets with nonzero values in the fields that must be zero. This is in spite of the fact that they are being sent in violation of the specifications in RFC 1058 and RFC 2453.	Select a value.
Graceful switchover	Configures graceful switchover for OSPF.	To disable graceful restart, select **Disable**. Type or select and edit the estimated time for the restart to finish, in seconds.
Authentication Type	Specifies the type of authentication for RIP route queries received on an interface. Options are: • None • MD5 • Simple	Select the authentication type. Enter the authentication key for MD5.
Policies tab		
Import Policy	Applies one or more policies to routes being imported into the local routing device from the neighbors.	Click **Add** to add an import policy. Click **Move up** or **Move down** to move the selected policy up or down the list of policies. Click **Remove** to remove an import policy.
Export Policy	Applies a policy to routes being exported to the neighbors.	Click **Add** to add an export policy. Click **Move up** or **Move down** to move the selected policy up or down the list of policies. Click **Remove** to remove an export policy.
Trace Options tab		
File Name	Specifies the name of the file to receive the output of the tracing operation.	Type or select and edit the name.
Number of Files	Specifies the maximum number of trace files.	Type or select and edit the name.
File Size	Specifies the maximum size for each trace file.	Type or select and edit the name.

Table 194: Edit RIP Global Settings *(continued)*

Field	Function	Your Action
World Readable	Specifies whether the trace file can be read by any user or not.	Select **True** to allow any user to read the file. Select **False** to disallow all users being able to read the file.
Flags	Specifies the tracing operation to perform.	Select a value from the list.

Related Documentation

- Monitoring RIP Routing Information on page 1460
- Layer 3 Protocols Supported on J-EX Series Switches on page 13

Configuring Static Routing (CLI Procedure)

Static routes are routes that are manually configured and entered into the routing table. Dynamic routes, in contrast, are learned by the J-EX Series switch and added to the routing table using a protocol such as OSPF or RIP.

The switch uses static routes:

- When the switch does not have a route to a destination that has a better (lower) *preference* value. The preference is an arbitrary value in the range from 0 through 255 that the software uses to rank routes received from different protocols, interfaces, or remote systems. The routing protocol process generally determines the active route by selecting the route with the lowest preference value. In the given range, **0** is the lowest and **255** is the highest.
- When the switch cannot determine the route to a destination.
- When the switch is forwarding unroutable packets.

To configure basic static route options using the CLI:

- To configure the switch's default gateway:

```
[edit]
user@switch# set routing–options static route 0.0.0.0/0 next-hop 10.0.1.1
```

- To configure a static route and specify the next address to be used when routing traffic to the static route:

```
[edit]
user@switch# set routing-options static route 20.0.0.0/24 next-hop 10.0.0.2.1
```

- To always keep the static route in the forwarding table:

```
[edit]
user@switch# set routing-options static route 20.0.0.0/24 retain
```

- To prevent the static route from being readvertised:

```
[edit]
user@switch# set routing-options static route 20.0.0.0/24 no-readvertise
```

- To remove inactive routes from the forwarding table:

```
[edit]
user@switch# set routing-options static route 20.0.0.0/24 active
```

Related Documentation

- Configuring Static Routing (J-Web Procedure) on page 1444
- Monitoring Routing Information on page 1461

Configuring Static Routing (J-Web Procedure)

You can use the J-Web interface to configure static routes for J-EX Series switches.

To configure static routes:

1. Select **Configure** > **Routing** > **Static Routing**. The Static Routing page displays details of the configured routes.

   NOTE: After you make changes to the configuration in this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options** > **Commit**. See "Using the Commit Options to Commit Configuration Changes (J-Web Procedure)" on page 334 for details about all commit options.

2. Click one:

   - **Add**—To configure a route. Enter information into the routing page as described in Table 195 on page 1445.
   - **Edit**—To modify an existing route. Enter information into the routing page as described in Table 195 on page 1445.
   - **Delete**—To delete an existing route.

Table 195: Static Routing Configuration Summary

Field	Function	Your Action
Default Route		
Default Route	Specifies the default gateway for the switch.	To specify an IPv4 address: 1. Select **IPv4**. 2. Type an IP address—for example, **10.10.10.10**. 3. Enter the subnet mask or address prefix. For example, 24 bits represents **255.255.255.0**. To specify an IPv6 address: 1. Select **IPv6**. 2. Type an IP address—for example, **2001:ab8:85a3::8a2e:370:7334**. 3. Enter the subnet mask or address prefix.
Static Routes		

Table 195: Static Routing Configuration Summary *(continued)*

Field	Function	Your Action
Nexthop	Specifies the next-hop address or addresses to be used when routing traffic to the static route.	To add an address: 1. Click **Add**. 2. In the IP address dialog, enter the IP address. NOTE: If a route has multiple next-hop addresses, traffic is routed across each address in round-robin fashion. 3. Click **OK**. To delete a next-hop address, select it from the list and click **Delete**.

Related Documentation

- Configuring Static Routing (CLI Procedure) on page 1444
- Monitoring Routing Information on page 1461
- Layer 3 Protocols Supported on J-EX Series Switches on page 13

Configuring Routing Policies (J-Web Procedure)

All routing protocols use the Junos OS routing table to store the routes that they learn and to determine which routes are advertised in the protocol packets. Routing policy allows you to control which routes the routing protocols store in and retrieve from the routing table on the routing device.

To configure routing policies for a J-EX Series switch using the J-Web interface:

1. Select **Configure** > **Routing** > **Policies**.

   NOTE: After you make changes to the configuration in this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options** > **Commit**. See "Using the Commit Options to Commit Configuration Changes (J-Web Procedure)" on page 334 for details about all commit options.

2. Click one:

   - **Global Options**—Configures global options for policies. Enter information into the configuration page as described in Table 196 on page 1447.
   - **Add**—Configures a new policy. Select **New** and specify a policy name. To add terms, enter information into the configuration page as described in Table 197 on page 1448. Select **Clone** to create a copy of an existing policy.

- **Edit**—Edits an existing policy. To modify an existing term, enter information into the configuration page as described in Table 197 on page 1448.
- **Term Up**—Moves a term up in the list.
- **Term Down**—Moves a term down in the list.
- **Delete**—Deletes the selected policy.
- **Test Policy**—Tests the policy. Use this option to check whether the policy produces the results that you expect.

Table 196: Policies Global Configuration Parameters

Field	Function	Your Action
Prefix List	Specifies a list of IPv4 address prefixes for use in a routing policy statement.	To add a prefix list: 1. Click **Add**. 2. Enter a name for the prefix list. 3. To add an IP address, click **Add**. 4. Enter the IP address and the subnet mask and click **OK**. 5. Click **OK**. To edit a prefix list, click **Edit**. Edit the settings and click **OK**. To delete a prefix list, select it and click **Delete**.
BGP Community	Specifies a BGP community.	To add a BGP community: 1. Click **Add**. 2. Enter a name for the community. 3. To add a community, click **Add**. 4. Enter the community ID and click **OK**. 5. Click **OK**. To edit a BGP community, click **Edit**. Edit the settings and click **OK**. To delete a BGP community, select it and click **Delete**.
AS Path	Specifies an AS path. This is applicable to BGP only.	To add an AS path: 1. Click **Add**. 2. Enter the AS path name. 3. Enter the regular expression and click **OK**. 4. Click **OK**. To edit an AS path, click **Edit**. Edit the settings and click **OK**. To delete an AS path, select it and click **Delete**.

Table 197: Terms Configuration Parameters

Field	Function	Your Action
Term Name	Specifies a term name.	Type or select and edit the name.
Source tab		
Family	Specifies an address family protocol.	Select a value from the list.
Routing Instance	Specifies a routing instance.	Select a value from the list.
RIB	Specifies the name of a routing table.	Select a value from the list
Preference	Specifies the individual preference value for the route.	Type or select and edit the value.
Metric	Specifies a metric value. You can specify up to four metric values.	Type or select and edit the value.
Interface	Specifies a name or IP address of one or more routing device interfaces. Do not use this qualifier with protocols that are not interface-specific, such as internal BGP (IBGP).	To add an interface, select **Add > Interface**. Select the interface from the list. To add an address, select **Add > Address**. Select the address from the list. To remove an interface, select it and click **Remove**.
Prefix List	Specifies a named list of IP addresses. You can specify an exact match with incoming routes.	Click **Add**. Select the prefix list from the list and click **OK**. To remove a prefix list, select it and click **Remove**.
Protocol	Specifies the name of the protocol from which the route was learned or to which the route is being advertised.	Click **Add** and select the protocol from the list. To remove a protocol, select it and click **Remove**.
Policy	Specifies the name of a policy to evaluate as a subroutine.	Click **Add**. Select the policy from the list. To remove a policy, select it and click **Remove**.
More	Specifies advanced configuration options for policies.	Click **More** for advanced configuration.
OSPF Area ID	Specifies the area identifier.	Type the IP address.
BGP Origin	Specifies the origin of the AS path information.	Select a value from the list.
Local Preference	Specifies the BGP local preference.	Type a value.

Table 197: Terms Configuration Parameters *(continued)*

Field	Function	Your Action
Route	Specifies the type of route.	Select **External**. Select the OSPF type from the list.
AS Path	Specifies the name of an AS path regular expression.	Click **Add**. Select the AS path from the list.
Community	Specifies the name of one or more communities.	Click **Add**. Select the community from the list.
Destination tab		
Family	Specifies an address family protocol.	Select a value from the list.
Routing Instance	Specifies a routing instance.	Select a value from the list.
RIB	Specifies the name of a routing table.	Select a value from the list.
Preference	Specifies the individual preference value for the route.	Type a value.
Metric	Specifies a metric value.	Type a value.
Interface	Specifies a name or IP address of one or more routing device interfaces. Do not use this qualifier with protocols that are not interface-specific, such as internal BGP (IBGP).	To add an interface, select **Add > Interface**. Select the interface from the list. To add an address, select **Add > Address**. Select the address from the list. To delete an interface, select it and click **Remove**.
Protocol	Specifies the name of the protocol from which the route was learned or to which the route is being advertised.	Click **Add** and select the protocol from the list. To delete a protocol, select it and click **Remove**.
Action tab		
Action	Specifies the action to take if the conditions match.	Select a value from the list.
Default Action	Specifies that any action that is intrinsic to the protocol is overridden. This action is also nonterminating, so that various policy terms can be evaluated before the policy is terminated.	Select a value from the list.
Next	Specifies the default control action if a match occurs, and there are no further terms in the current routing policy.	Select a value from the list.
Priority	Specifies a priority for prefixes included in an OSPF import policy. Prefixes learned through OSPF are installed in the routing table based on the priority assigned to the prefixes.	Select a value from the list.
BGP Origin	Specifies the BGP origin attribute.	Select a value from the list.

Table 197: Terms Configuration Parameters *(continued)*

Field	Function	Your Action
AS Path Prepend	Affixes an AS number at the beginning of the AS path. The AS numbers are added after the local AS number has been added to the path. This action adds an AS number to AS sequences only, not to AS sets. If the existing AS path begins with a confederation sequence or set, the affixed AS number is placed within a confederation sequence. Otherwise, the affixed AS number is placed with a nonconfederation sequence.	Enter a value.
AS Path Expand	Extracts the last AS number in the existing AS path and affixes that AS number to the beginning of the AS path *n* times, where *n* is a number from 1 through 32. The AS number is added before the local AS number has been added to the path. This action adds AS numbers to AS sequences only, not to AS sets. If the existing AS path begins with a confederation sequence or set, the affixed AS numbers are placed within a confederation sequence. Otherwise, the affixed AS numbers are placed within a nonconfederation sequence. This option is typically used in non-IBGP export policies.	Select the type and type a value.
Load Balance Per Packet	Specifies that all next-hop addresses in the forwarding table must be installed and have the forwarding table perform per-packet load balancing. This policy action allows you to optimize VPLS traffic flows across multiple paths.	Select the check box to enable the option.
Tag	Specifies the tag value. The tag action sets the 32-bit tag field in OSPF external link-state advertisement (LSA) packets.	Select the action and type a value.
Metric	Changes the metric (MED) value by the specified negative or positive offset. This action is useful only in an external BGP (EBGP) export policy.	Select the action and type a value.
Route	Specifies whether the route is external.	Select the **External** check box to enable the option, and select the OSPF type.
Preference	Specifies the preference value.	Select the preference action and type a value.
Local Preference	Specifies the BGP local preference attribute.	Select the action and type a value.
Class of Service	Specifies and applies the class-of-service parameters to routes installed into the routing table. • Source class The value entered here maintains the packet counts for a route passing through your network, based on the source address. • Destination class The value entered here maintains packet counts for a route passing through your network, based on the destination address in the packet. • Forwarding class	Type the source class. Type the destination class. Type the forwarding class.

Related Documentation

- Configuring BGP Sessions (J-Web Procedure) on page 1431
- Configuring an OSPF Network (J-Web Procedure) on page 1435
- Configuring a RIP Network (J-Web Procedure) on page 1439
- Configuring Static Routing (J-Web Procedure) on page 1444
- Layer 3 Protocols Supported on J-EX Series Switches on page 13

Configuring Distributed Periodic Packet Management on a J-EX Series Switch (CLI Procedure)

Periodic packet management (PPM) is responsible for processing a variety of time-sensitive periodic tasks so that other processes on the J-EX Series switch can more optimally direct their resources.

The responsibility for PPM processing on the switch is distributed between the Routing Engine and either the access interfaces (on J-EX4200 switches) or the line cards (on J-EX8200 switches) for all protocols that use PPM by default. This distributed model provides a faster response time for protocols that use PPM than the response time provided by the nondistributed model.

If distributed PPM is disabled, the PPM process runs on the Routing Engine only.

Distributed PPM can be disabled for all protocols that use PPM or for a single protocol that uses PPM.

BEST PRACTICE: We recommend that, generally, you disable distributed PPM only if Dell Support advises you to do so (see “Requesting Technical Support” on page lxxi). You should disable distributed PPM only if you have a compelling reason to disable it.

This topic describes:

- Disabling or Enabling Distributed Periodic Packet Management Globally on page 1451
- Disabling or Enabling Distributed Periodic Packet Management for Link Aggregation Control Protocol (LACP) Packets on page 1452

Disabling or Enabling Distributed Periodic Packet Management Globally

Distributed PPM is enabled by default. Disable distributed PPM if you need to move all PPM processing to the Routing Engine. Enable distributed PPM if it was previously disabled and you need to run distributed PPM.

To disable distributed PPM:

```
[edit routing-options]
user@switch# set ppm no-delegate-processing
```

To enable distributed PPM if it was previously disabled:

```
[edit routing-options]
```

```
user@switch# delete ppm no-delegate-processing
```

Disabling or Enabling Distributed Periodic Packet Management for Link Aggregation Control Protocol (LACP) Packets

Distributed PPM is enabled by default. Disable distributed PPM for only LACP packets if you need to move all PPM processing for LACP packets to the Routing Engine.

To disable distributed PPM for LACP packets:

```
[edit protocols]
user@switch# set lacp ppm centralized
```

To enable distributed PPM for LACP packets if it was previously disabled:

```
[edit protocols]
user@switch# delete lacp ppm centralized
```

Related Documentation

- Understanding Distributed Periodic Packet Management on J-EX Series Switches on page 1424
- Understanding Aggregated Ethernet Interfaces and LACP on page 867

Configuring VRRP for IPv6 (CLI Procedure)

By configuring the Virtual Router Redundancy Protocol (VRRP) on J-EX Series switches, you can enable hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts. You can configure VRRP for IPv6 on Gigabit Ethernet, 10-Gigabit Ethernet, and logical interfaces.

To configure VRRP for IPv6:

1. Configure VRRP group support on interfaces:

   ```
   [edit interfaces interface-name unit logical-unit-number family inet6
   address address]
   user@switch# set vrrp-inet6-group group-id priority number virtual-inet6-address
   address virtual-link-local-address ipv6-address
   ```

   You must explicitly define a virtual link local address for each VRRP for IPv6 group. Otherwise, when you attempt to commit the configuration, the commit request fails. The virtual link local address must be on the same subnet as the physical interface address.

2. If you want to configure the priority order in which this switch functioning as a backup router becomes the master router if the master router becomes nonoperational, configure a priority for this switch:

   ```
   [edit interfaces interface-name unit logical-unit-number family inet6
   address address vrrp-inet6-group group-id]
   user@switch# set priority number
   ```

3. Specify the interval in milliseconds in which the master router sends advertisement packets to the members of the VRRP group:

   ```
   [edit interfaces interface-name unit logical-unit-number family inet6
   address address vrrp-inet6-group group-id]
   ```

```
user@switch# set inet6-advertise-interval milliseconds
```

4. By default, a higher-priority backup router preempts a lower-priority master router.

   - To explicitly enable the master router to be preempted:

     ```
     [edit interfaces interface-name unit logical-unit-number family inet6
     address address vrrp-inet6-group group-id]
     user@switch# set preempt
     ```

   - To prohibit a higher-priority backup router from preempting a lower priority master router:

     ```
     [edit interfaces interface-name unit logical-unit-number family inet6
     address address vrrp-inet6-group group-id]
     user@switch# set no-preempt
     ```

Related Documentation

- **show vrrp on page 2036**
- Understanding VRRP on J-EX Series Switches on page 1425

Using IPsec to Secure OSPFv3 Networks (CLI Procedure)

OSPF version 3 (OSPFv3) does not have a built-in authentication method and relies on IP Security (IPsec) to provide this functionality. You can use IPsec to secure OSPFv3 interfaces on J-EX Series switches.

This topic includes:

Configuring Security Associations

When you configure a security association (SA), include your choices for authentication, encryption, direction, mode, protocol, and security parameter index (SPI).

To configure a security association:

1. Specify a name for the security association:

   ```
   [edit security ipsec]
   user@switch# set security-association sa-name
   ```

2. Specify the mode of the security association:

   ```
   [edit security ipsec security-association sa-name]
   user@switch# set mode transport
   ```

3. Specify the type of security association:

   ```
   [edit security ipsec security-association sa-name]
   user@switch# set type manual
   ```

4. Specify the direction of the security association:

   ```
   [edit security ipsec security-association sa-name]
   ```

```
user@switch# set direction bidirectional
```

5. Specify the value of the security parameter index:

```
[edit security ipsec security-association sa-name]
user@switch# set spi spi-value
```

6. Specify the type of authentication to be used:

```
[edit security ipsec security-association sa-name]
user@switch# set authentication algorithm type
```

7. Specify the encryption algorithm and key:

```
[edit security ipsec security-association sa-name]
user@switch# set encryption algorithm algorithm key type
```

Securing OPSFv3 Networks

You can secure the OSPFv3 network by applying the SA to the OSPFv3 configuration.

To secure the OSPFv3 network:

```
[edit protocols ospf3 area area-number interface interface-name]
user@switch# set ipsec-sa sa-name
```

Related Documentation

- Understanding IPsec Authentication for OSPF Packets on J-EX Series Switches on page 1428
- Configuring an OSPF Network (J-Web Procedure) on page 1435
- For details on these configuration statements, see the *Junos OS System Basics Configuration Guide* at http://www.juniper.net/techpubs/software/junos/index.html.

CHAPTER 72

Verifying Layer 3 Protocols Configuration

Monitoring BGP Routing Information

Purpose Use the monitoring functionality to monitor BGP routing information on the routing device.

Action To view BGP routing information in the J-Web interface, select **Monitor>Routing>BGP Information**.

To view BGP routing information in the CLI, enter the following commands:

- **show bgp summary**
- **show bgp neighbor**

Meaning Table 198 on page 1455 summarizes key output fields in the BGP routing display in the J-Web interface.

Table 198: Summary of Key BGP Routing Output Fields

Field	Values	Additional Information
BGP Peer Summary		
Total Groups	Number of BGP groups.	
Total Peers	Number of BGP peers.	
Down Peers	Number of unavailable BGP peers.	
Unconfigured Peers	Address of each BGP peer.	
RIB Summary tab		
RIB Name	Name of the RIB group.	

Table 198: Summary of Key BGP Routing Output Fields *(continued)*

Field	Values	Additional Information
Total Prefixes	Total number of prefixes from the peer, both active and inactive, that are in the routing table.	
Active Prefixes	Number of prefixes received from the EBGP peers that are active in the routing table.	
Suppressed Prefixes	Number of routes received from EBGP peers currently inactive because of damping or other reasons.	
History Prefixes	History of the routes received or suppressed.	
Dumped Prefixes	Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.	
Pending Prefixes	Number of pending routes.	
State	Status of the graceful restart process for this routing table: BGP restart is complete, BGP restart in progress, VPN restart in progress, or VPN restart is complete.	
BGP Neighbors		
Details	Click this button to view the selected BGP neighbor details.	
Peer Address	Address of the BGP neighbor.	
Autonomous System	AS number of the peer.	

Table 198: Summary of Key BGP Routing Output Fields *(continued)*

Field	Values	Additional Information
Peer State	Current state of the BGP session: • **Active**—BGP is initiating a TCP connection in an attempt to connect to a peer. If the connection is successful, BGP sends an open message. • **Connect**—BGP is waiting for the TCP connection to become complete. • **Established**—The BGP session has been established, and the peers are exchanging BGP update messages. • **Idle**—This is the first stage of a connection. BGP is waiting for a Start event. • **OpenConfirm**—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message. • **OpenSent**—BGP has sent an open message and is waiting to receive an open message from the peer.	Generally, the most common states are **Active**, which indicates a problem establishing the BGP connection, and **Established**, which indicates a successful session setup. The other states are transition states, and BGP sessions normally do not stay in those states for extended periods of time.
Elapsed Time	Elapsed time since the peering session was last reset.	
Description	Description of the BGP session.	

Related Documentation

- Configuring BGP Sessions (J-Web Procedure) on page 1431
- Layer 3 Protocols Supported on J-EX Series Switches on page 13

Monitoring OSPF Routing Information

Purpose Use the monitoring functionality to monitor OSPF routing information on routing devices.

Action To view OSPF routing information in the J-Web interface, select **Monitor > Routing > OSPF Information**.

To view OSPF routing information in the CLI, enter the following CLI commands:

- **show ospf neighbor**
- **show ospf interface**
- **show ospf statistics**

Meaning Table 199 on page 1458 summarizes key output fields in the OSPF routing display in the J-Web interface.

Table 199: Summary of Key OSPF Routing Output Fields

Field	Values	Additional Information
OSPF Interfaces		
Interface	Name of the interface running OSPF.	
State	State of the interface: **BDR**, **Down**, **DR**, **DRother**, **Loop**, **PtToPt**, or **Waiting**.	The **Down** state, indicating that the interface is not functioning, and **PtToPt** state, indicating that a point-to-point connection has been established, are the most common states.
Area	Number of the area that the interface is in.	
DR ID	Address of the area's designated device.	
BDR ID	Address of the area's backup designated device.	
Neighbors	Number of neighbors on this interface.	
Adjacency Count	Number of devices in the area using the same area identifier.	
Stub Type	The areas into which OSPF does not flood AS external advertisements	
Passive Mode	In this mode the interface is present on the network but does not transmit or receive packets.	
Authentication Type	The authentication scheme for the backbone or area.	
Interface Address	The IP address of the interface.	
Address Mask	The subnet mask or address prefix.	
MTU	The maximum transmission unit size.	
Interface Cost	The path cost used to calculate the root path cost from any given LAN segment is determined by the total cost of each link in the path.	
Hello Interval	How often the routing device sends hello packets out of the interface.	
Dead Interval	The interval during which the routing device receives no hello packets from the neighbor.	
Retransmit Interval	The interval for which the routing device waits to receive a link-state acknowledgment packet before retransmitting link-state advertisements to an interface's neighbors.	

Table 199: Summary of Key OSPF Routing Output Fields *(continued)*

Field	Values	Additional Information
OSPF Statistics		
Packets tab		
Sent	Displays the total number of packets sent.	
Received	Displays the total number of packets received.	
Details tab		
Flood Queue Depth	Number of entries in the extended queue.	
Total Retransmits	Number of retransmission entries enqueued.	
Total Database Summaries	Total number of database description packets.	
OSPF Neighbors		
Address	Address of the neighbor.	
Interface	Interface through which the neighbor is reachable.	
State	State of the neighbor: **Attempt**, **Down**, **Exchange**, **ExStart**, **Full**, **Init**, **Loading**, or **2way**.	Generally, only the **Down** state, indicating a failed OSPF adjacency, and the **Full** state, indicating a functional adjacency, are maintained for more than a few seconds. The other states are transitional states that a neighbor is in only briefly while an OSPF adjacency is being established.
ID	ID of the neighbor.	
Priority	Priority of the neighbor to become the designated router.	
Activity Time	The activity time.	
Area	Area that the neighbor is in.	
Options	Option bits received in the hello packets from the neighbor.	
DR Address	Address of the designated router.	
BDR Address	Address of the backup designated router.	
Uptime	Length of time since the neighbor came up.	

Table 199: Summary of Key OSPF Routing Output Fields *(continued)*

Field	Values	Additional Information
Adjacency	Length of time since the adjacency with the neighbor was established.	

Related Documentation

- Configuring an OSPF Network (J-Web Procedure) on page 1435
- Layer 3 Protocols Supported on J-EX Series Switches on page 13

Monitoring RIP Routing Information

Purpose Use the monitoring functionality to monitor RIP routing on routing devices.

Action To view RIP routing information in the J-Web interface, select **Monitor > Routing > RIP Information**.

To view RIP routing information in the CLI, enter the following CLI commands:

- **show rip statistics**
- **show rip neighbor**

Meaning Table 200 on page 1460 summarizes key output fields in the RIP routing display in the J-Web interface.

Table 200: Summary of Key RIP Routing Output Fields

Field	Values	Additional Information
RIP Statistics		
Protocol Name	The RIP protocol name.	
Port number	The port on which RIP is enabled.	
Hold down time	The interval during which routes are neither advertised nor updated.	
Global routes learned	Number of RIP routes learned on the logical interface.	
Global routes held down	Number of RIP routes that are not advertised or updated during the hold-down interval.	
Global request dropped	Number of requests dropped.	
Global responses dropped	Number of responses dropped.	

Table 200: Summary of Key RIP Routing Output Fields *(continued)*

Field	Values	Additional Information
RIP Neighbors		
Neighbor	Name of the RIP neighbor.	This value is the name of the interface on which RIP is enabled. Click the name to see the details for this neighbor.
State	State of the RIP connection: **Up** or **Dn** (Down).	
Source Address	Local source address.	This value is the configured address of the interface on which RIP is enabled.
Destination Address	Destination address.	This value is the configured address of the immediate RIP adjacency.
Send Mode	The mode of sending RIP messages.	
Receive Mode	The mode in which messages are received.	
In Metric	Value of the incoming metric configured for the RIP neighbor.	

Related Documentation

- Configuring a RIP Network (J-Web Procedure) on page 1439
- Layer 3 Protocols Supported on J-EX Series Switches on page 13

Monitoring Routing Information

Purpose Use the monitoring functionality to view the **inet.0** routing table on the routing device.

Action To view the routing tables in the J-Web interface, select **Monitor > Routing > Route Information**. Apply a filter or a combination of filters to view messages. You can use filters to display relevant events.

To view the routing table in the CLI, enter the following commands in the CLI interface:

- **show route terse**
- **show route detail**

Meaning Table 201 on page 1462 describes the different filters, their functions, and the associated actions.

Table 202 on page 1462 summarizes key output fields in the routing information display.

Table 201: Filtering Route Messages

Field	Function	Your Action
Destination Address	Specifies the destination address of the route.	Enter the destination address.
Protocol	Specifies the protocol from which the route was learned.	Enter the protocol name.
Next hop address	Specifies the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.	Enter the next hop address.
Receive protocol	Specifies the dynamic routing protocol using which the routing information was received through a particular neighbor.	Enter the routing protocol.
Best route	Specifies only the best route available.	Select the view details of the best route.
Inactive routes	Specifies the inactive routes.	Select the view details of inactive routes.
Exact route	Specifies the exact route.	Select the view details of the exact route.
Hidden routes	Specifies the hidden routes.	Select the view details of hidden routes.
Search	Applies the specified filter and displays the matching messages.	To apply the filter and display messages, click **Search**.

Table 202: Summary of Key Routing Information Output Fields

Field	Values	Additional Information
Static Route Addresses	The list of static route addresses.	
Protocol	Protocol from which the route was learned: **Static**, **Direct**, **Local**, or the name of a particular protocol.	
Preference	The preference is the individual preference value for the route.	The route preference is used as one of the route selection criteria.

Table 202: Summary of Key Routing Information Output Fields *(continued)*

Field	Values	Additional Information
Next-Hop	Network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.	If a next hop is listed as **Discard**, all traffic with that destination address is discarded rather than routed. This value generally means that the route is a static route for which the **discard** attribute has been set. If a next hop is listed as **Reject**, all traffic with that destination address is rejected. This value generally means that the address is unreachable. For example, if the address is a configured interface address and the interface is unavailable, traffic bound for that address is rejected. If a next hop is listed as **Local**, the destination is an address on the host (either the loopback address or Ethernet management port 0 address, for example).
Age	How long the route has been active.	
State	Flags for this route.	There are many possible flags.
AS Path	AS path through which the route was learned. The letters of the AS path indicate the path origin: • **I**—IGP. • **E**—EGP. • **?**—Incomplete. Typically, the AS path was aggregated.	

Related Documentation

- Configuring Static Routing (J-Web Procedure) on page 1444
- Configuring Static Routing (CLI Procedure) on page 1444
- Layer 3 Protocols Supported on J-EX Series Switches on page 13

CHAPTER 73

Configuration Statements for Layer 3 Protocols

accept-remote-nexthop

Syntax accept-remote-nexthop;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp group *group-name*],
[edit logical-systems *logical-system-name* protocols bgp group *group-name* neighbor *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address*],
[edit protocols bgp],
[edit protocols bgp group *group-name*],
[edit protocols bgp group *group-name* neighbor *address*],
[edit routing-instances *routing-instance-name* protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify that a single-hop EBGP peer accept a remote next hop with which it does not share a common subnet. Configure a separate import policy on the EBGP peer to specify the remote next hop. You cannot configure the **multihop** statement at the same time.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **multipath on page 1623**
- Configuring Single-Hop EBGP Peers to Accept Remote Next Hops
- Applying Policies to BGP Routes

active

Syntax (active | passive);

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options (**aggregate** | **generate** | **static**) (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* (**aggregate** | **generate** | **static**) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options (**aggregate** | **generate** | **static**) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* (**aggregate** | **generate** | **static**) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options (**aggregate** | **generate** | **static**) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* (**aggregate** | **generate** | **static**) (defaults | route)],
[edit routing-options (**aggregate** | **generate** | **static**) (defaults | route)],
[edit routing-options rib *routing-table-name* (**aggregate** | **generate** | **static**) (defaults | route)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure whether static, aggregate, or generated routes are removed from the routing and forwarding tables when they become inactive. Routes that have been configured to remain continually installed in the routing and forwarding tables are marked with **reject** next hops when they are inactive.

- **active**—Remove a route from the routing and forwarding tables when it becomes inactive.
- **passive**—Have a route remain continually installed in the routing and forwarding tables even when it becomes inactive.

Default **active**

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Static Routes
- Configuring Aggregate Routes
- Configuring Generated Routes

advertise-external

Syntax

```
advertise-external {
  conditional;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols bgp group *group-name*],
[edit logical-systems *logical-system-name* protocols bgp group *group-name* neighbor *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address*],
[edit protocols bgp group *group-name*],
[edit protocols bgp group *group-name* neighbor *address*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *neighbor-address*]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Have BGP advertise the best external route into an IBGP mesh group, a route reflector cluster, or an AS confederation even if the best route is an internal route.

Options

conditonal—(Optional) Advertise the best external path only if the route selection process reaches the point where the multiple exit discriminator (MED) metric is evaluated. As a result, an external path with an AS path worse than that of the active path is not advertised.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **advertise-inactive on page 1468**
- Applying Policies to BGP Routes

advertise-inactive

Syntax advertise-inactive;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name* neighbor *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],
[edit protocols bgp],
[edit protocols bgp **group** *group-name*],
[edit protocols bgp group *group-name* **neighbor** *address*],
[edit routing-instances *routing-instance-name* protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Have BGP advertise the best route even if the routing table did not select it to be an active route.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Applying Policies to BGP Routes

advertise-peer-as

Syntax advertise-peer-as;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name* neighbor *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],
[edit protocols bgp],
[edit protocols bgp **group** *group-name*],
[edit protocols bgp group *group-name* **neighbor** *address*],
[edit routing-instances *routing-instance-name* protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable the default behavior of suppressing AS routes.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Applying Policies to BGP Routes

aggregate

Syntax

```
aggregate {
  defaults {
    ... aggregate-options ...
  }
  route destination-prefix {
    policy policy-name;
    ... aggregate-options ...
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options **rib** *routing-table-name*],
[edit logical-systems *logical-system-name* routing-options],
[edit logical-systems *logical-system-name* routing-options **rib** *routing-table-name*],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options **rib** *routing-table-name*],
[edit routing-options],
[edit routing-options **rib** *routing-table-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure aggregate routes.

Options ***aggregate-options***—Additional information about aggregate routes that is included with the route when it is installed in the routing table. Specify zero or more of the following options in ***aggregate-options***. Each option is explained separately.

- **(active | passive);**
- **as-path <***as-path***> <origin (egp | igp | incomplete)> <atomic-aggregate> <aggregator** *as-number in-address***>**;
- **(brief | full)**;
- **community [** *community-ids* **]**;
- **discard**;
- **(metric | metric2 | metric3 | metric4)** *value* **<type** *type***>**;
- **(preference | preference2 | color | color2)** *preference* **<type** *type***>**;
- **tag** *string*;

defaults—Specify global aggregate route options. These options only set default attributes inherited by all newly created aggregate routes. These are treated as global defaults and apply to all the aggregate routes you configure in the **aggregate** statement. This part of the **aggregate** statement is optional.

route ***destination-prefix***—Configure a nondefault aggregate route:

- **default**—For the default route to the destination. This is equivalent to specifying an IP address of **0.0.0.0/0**.
- ***destination-prefix/prefix-length*—*destination-prefix*** is the network portion of the IP address, and ***prefix-length*** is the destination prefix length.

The **policy** statement is explained separately.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Aggregate Routes

aggregate-label

Syntax

```
aggregate-label {
  community community-name;
}
```

Hierarchy Level
[edit logical-systems *logical-system-name* protocols bgp family inet labeled-unicast],
[edit logical-systems *logical-system-name* protocols bgp family inet-vpn labeled-unicast],
[edit protocols bgp family inet labeled-unicast],
[edit protocols bgp family inet-vpn labeled-unicast],
[edit protocols bgp family inet6 labeled-unicast]

Release Information
Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Enable aggregate labels for VPN traffic.

Options
community *community-name*—Specify the name of the community to which to apply the aggregate label.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Aggregate Labels for VPNs

allow

Syntax allow (all | [*network/mask-length*]);

Hierarchy Level [edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Implicitly configure BGP peers, allowing peer connections from any of the specified networks or hosts. To configure multiple BGP peers, configure one or more networks and hosts within a single **allow** statement or include multiple **allow** statements.

Options **all**—Allow all addresses, which is equivalent to **0.0.0.0/0** (or **::/0**).

network/mask-length—IPv6 or IPv4 network number of a single address or a range of allowable addresses for BGP peers, followed by the number of significant bits in the subnet mask.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **neighbor on page 1624**
- Minimum BGP Configuration
- Configuring BGP Groups and Peers

any-sender

Syntax any-sender;

Hierarchy Level [edit logical-systems *logical-system-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit protocols rip group *group-name* neighbor *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable strict sender address checks.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Disabling Strict Address Checking for RIP Messages

area

Syntax area *area-id*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols (**ospf** | **ospf3**)],
[edit logical-systems *logical-system-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit protocols (**ospf** | **ospf3**)],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)],
[edit routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the area identifier for this routing device to use when participating in OSPF routing. All routing devices in an area must use the same area identifier to establish adjacencies.

Specify multiple **area** statements to configure the routing device as an area border router. An area border router does not automatically summarize routes between areas; use the **area-range** statement to configure route summarization. By definition, an area border router must be connected to the backbone area either through a physical link or through a virtual link. To create a virtual link, include the **virtual-link** statement.

To specify that the routing device is directly connected to the OSPF and OSPFv3 backbone, include the **area 0.0.0.0** statement.

All routing devices on the backbone must be contiguous. If they are not, use the **virtual-link** statement to create the appearance of connectivity to the backbone.

Options *area-id*—Area identifier. The identifier can be up to 32 bits. It is common to specify the area number as a simple integer or an IP address. Area number **0.0.0.0** is reserved for the OSPF and OSPFv3 backbone area.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **virtual-link on page 1742**
- Configuring OSPF Areas
- Configuring Multiple Address Families for OSPFv3

area-range

Syntax area-range *network/mask-length* <exact> <override-metric *metric*> <restrict>;

Hierarchy Level [edit logical-systems *logical-system-name* protocols (ospf | ospf3) **area** *area-id*],
[edit logical-systems *logical-system-name* protocols (ospf | ospf3) area *area-id* **nssa**],
[edit logical-systems *logical-system-name* **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf | ospf3) **area** *area-id*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* **nssa**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) **area** *area-id*],
[edit protocols (ospf | ospf3) **area** *area-id*],
[edit protocols (ospf | ospf3) area *area-id* **nssa**],
[edit protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) **area** *area-id*],
[edit routing-instances *routing-instance-name* protocols (ospf | ospf3) **area** *area-id*],
[edit routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* **nssa**],
[edit routing-instances *routing-instance-name* **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) **area** *area-id*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description (Area border routers only) For an area, summarize a range of IP addresses when sending summary link advertisements (within an area). To summarize multiple ranges, include multiple **area-range** statements.

For a not-so-stubby area (NSSA), summarize a range of IP addresses when sending NSSA link-state advertisements. The specified prefixes are used to aggregate external routes learned within the area when the routes are advertised to other areas. To specify multiple prefixes, include multiple **area-range** statements. All external routes learned within the area that do not fall into one of the prefixes are advertised individually to other areas.

Default By default, area border routers do not summarize routes being sent from one area to other areas, but rather send all routes explicitly.

Options

exact—(Optional) Summarization of a route is advertised only when an exact match is made with the configured summary range.

mask-length—Number of significant bits in the network mask.

network—IP address. You can specify one or more IP addresses.

override-metric *metric*—(Optional) Override the metric for the IP address range and configure a specific metric value.

restrict—(Optional) Do not advertise the configured summary. This hides all routes that are contained within the summary, effectively creating a route filter.

Range: 1 through 16,777,215

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Summarizing Ranges of Routes in OSPF Link-State Advertisements

as-override

Syntax
as-override;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name* neighbor *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],
[edit protocols bgp **group** *group-name*],
[edit protocols bgp group *group-name* **neighbor** *address*],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*]

Release Information
Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Compare the AS path of an incoming advertised route with the AS number of the BGP peer under the group and replace all occurrences of the peer AS number in the AS path with its own AS number before advertising the route to the peer.

NOTE: The **as-override** statement is specific to a particular BGP group. This statement does not affect peers from the same remote AS configured in different groups.

Enabling the AS override feature allows routes originating from an AS to be accepted by a router residing in the same AS. Without AS override enabled, the routing device refuses the route advertisement once the AS path shows that the route originated from its own AS. This is done by default to prevent route loops. The **as-override** statement overrides this default behavior.

Note that enabling the AS override feature may result in routing loops. Use this feature only for specific applications that require this type of behavior, and in situations with strict network control. One application is the IGP protocol between the provider edge routing device and the customer edge routing device in a virtual private network. For more information, see the *Junos OS MPLS Applications Configuration Guide*.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring BGP Groups and Peers

as-path

Syntax as-path <*as-path*> <aggregator *as-number ip-address*> <atomic-aggregate> <origin (egp | igp | incomplete)>;

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options (**aggregate | generate | static**) (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* (**aggregate | generate | static**) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options (**aggregate | generate | static**) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* (**aggregate | generate | static**) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options (**aggregate | generate | static**) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* (**aggregate | generate | static**) (defaults | route)],
[edit routing-options (**aggregate | generate | static**) (defaults | route)],
[edit routing-options rib *routing-table-name* (**aggregate | generate | static**) (defaults | route)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Associate BGP autonomous system (AS) path information with a static, aggregate, or generated route.

The numeric range for the AS number provides BGP support for 2-byte AS numbers and 4-byte AS numbers. .

Options

aggregator—(Optional) Attach the BGP **aggregator** path attribute to the aggregate route. You must specify the last AS number that formed the aggregate route (encoded as two octets) for ***as-number***, followed by the IP address of the BGP system that formed the aggregate route for ***in-address***.

as-path—(Optional) AS path to include with the route. It can include a combination of individual AS path numbers and AS sets. Enclose sets in brackets ([ ]). The first AS number in the path represents the AS immediately adjacent to the local AS. Each subsequent number represents an AS that is progressively farther from the local AS, heading toward the origin of the path. You cannot specify a regular expression for ***as-path***; you must use a full, valid AS path.

atomic-aggregate—(Optional) Attach the BGP **atomic-aggregate** path attribute to the aggregate route. This path attribute indicates that the local system selected a less specific route instead of a more specific route.

origin egp—(Optional) BGP origin attribute that indicates that the path information originated in another AS.

origin igp—(Optional) BGP origin attribute that indicates that the path information originated within the local AS.

origin incomplete—(Optional) BGP origin attribute that indicates that the path information was learned by some other means.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Static Routes
- Configuring Aggregate Routes
- Configuring Generated Routes

asm-override-ssm

Syntax
asm-override-ssm;

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information
Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Enable the routing device to accept any-source multicast join messages (*,G) for group addresses that are within the default or configured range of source-specific multicast groups.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Source-Specific Multicast Groups

authentication-algorithm

Syntax authentication-algorithm *algorithm*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name* neighbor *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],
[edit protocols bgp],
[edit protocols bgp **group** *group-name*],
[edit protocols bgp group *group-name* **neighbor** *address*],
[edit routing-instances *routing-instance-name* protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure an authentication algorithm type.

Options ***algorithm***—Type of authentication algorithm. Specify **md5**, **hmac-sha-1-96**, or **aes-128-cmac-96** as the algorithm type.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Authentication for BGP

authentication-key

Syntax authentication-key *key*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name* **neighbor** *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],
[edit protocols bgp],
[edit protocols bgp **group** *group-name*],
[edit protocols bgp group *group-name* **neighbor** *address*],
[edit routing-instances *routing-instance-name* protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure an MD5 authentication key (password). Neighboring routing devices use the same password to verify the authenticity of BGP packets sent from this system.

Options ***key***—Authentication password. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (“ ”).

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Authentication for BGP

authentication-key

Syntax authentication-key *key*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols isislevel *level-number*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis **level** *level-number*],
[edit protocols isis **level** *level-number*],
[edit routing-instances *routing-instance-name* protocols isis **level** *level-number*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Authentication key (password). Neighboring routing devices use the password to verify the authenticity of packets sent from this interface. For the key to work, you also must include the **authentication-type** statement.

All routing devices must use the same password. If you are using the Junos OS IS-IS software with another implementation of IS-IS, the other implementation must be configured to use the same password for the domain, the area, and all interfaces adjacent to the Dell PowerConnect J-Series routing device.

Default If you do not include this statement and the **authentication-type** statement, IS-IS authentication is disabled.

Options *key*—Authentication password. The password can be up to 1024 characters long. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (“ ”).

CAUTION: A simple password for authentication is truncated if it exceeds 254 characters.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring IS-IS Authentication

authentication-key

Syntax authentication-key *password*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols **rip**],
[edit logical-systems *logical-system-name* protocols rip group *group-name* **neighbor** *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **rip**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip group *group-name* **neighbor** *neighbor-name*],
[edit protocols **rip**],
[edit protocols rip group *group-name* **neighbor** *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols **rip**],
[edit routing-instances *routing-instance-name* protocols rip group *group-name* **neighbor** *neighbor-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Require authentication for RIP route queries received on an interface.

Options ***password***—Authentication password. If the password does not match, the packet is rejected. The password can be from 1 through 16 contiguous characters long and can include any ASCII strings.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Authentication for RIP

authentication-key-chain

Syntax authentication-key-chain *key-chain*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name* **neighbor** *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],
[edit protocols bgp],
[edit protocols bgp **group** *group-name*],
[edit protocols bgp group *group-name* **neighbor** *address*],
[edit routing-instances *routing-instance-name* protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply and enable an authentication keychain to the routing device.

Options ***key-chain***—Authentication keychain name. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (“ ”).

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Authentication for BGP

authentication-key-chains

Syntax

```
authentication-key-chains {
  key-chain key-chain-name {
    description text-string;
    key key {
      secret secret-data;
      start-time yyyy-mm-dd.hh:mm:ss;
    }
    tolerance seconds;
  }
}
```

Hierarchy Level [edit security]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure authentication key updates for the Border Gateway Protocol (BGP), the Label Distribution Protocol (LDP) routing protocols, and the Bidirectional Forwarding Detection (BFD) protocol. When the **authentication-key-chains** statement is configured at the **[edit security]** hierarchy level, and is associated with the BGP and LDP protocols at the **[edit protocols]** hierarchy level or with the BFD protocol using the **bfd-liveness-detection** statement, authentication key updates can occur without interrupting routing and signaling protocols such as Open Shortest Path First (OSPF), and Resource Reservation Setup Protocol (RSVP).

Options

key-chain *key-chain-name*—Keychain name. This name is configured at the **[edit protocols bgp]** or the **[edit protocols ldp]** hierarchy level to associate unique **authentication key-chain** attributes with each protocol as specified using the following options:

- **description *text-string***—A text string of the **authentication-key-chain**. Put the text string in quotes (“text description”).
- **key *key***—Each key within a keychain is identified by a unique integer value.

  Range: 0 through **63**

  - **secret *secret-data***—Each key must specify a secret in encrypted text or plain text format. The secret always appears in encrypted format.
  - **start-time *yyyy–mm-dd.hh:mm:ss***—Start times are specified in UTC (Coordinated Universal Time), and must be unique within the keychain.
- **tolerance *seconds***—Specify the clock skew tolerance, in seconds.

  Range: 0 through **999999999**

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols

- Configuring BFD Authentication for Static Routes

authentication-type

Syntax authentication-type *authentication*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols isis **level** *level-number*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis **level** *level-number*],
[edit protocols isis **level** *level-number*],
[edit routing-instances *routing-instance-name* protocols isis **level** *level-number*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable authentication and specify the authentication scheme for IS-IS. If you enable authentication, you must specify a password by including the **authentication-key** statement.

Default If you do not include this statement and the **authentication-key** statement, IS-IS authentication is disabled.

Options ***authentication***—Authentication scheme:

- **md5**—Use HMAC authentication in combination with MD5. HMAC-MD5 authentication is defined in RFC 2104, *HMAC: Keyed-Hashing for Message Authentication*.
- **simple**—Use a simple password for authentication. The password is included in the transmitted packet, making this method of authentication relatively insecure. We recommend that you *not* use this authentication method.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **authentication-key on page 1481**
- **no-authentication-check on page 1631**
- Configuring IS-IS Authentication

authentication-type

Syntax authentication-type *type*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols **rip**],
[edit logical-systems *logical-system-name* protocols rip group *group-name* **neighbor** *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **rip**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip group *group-name* **neighbor** *neighbor-name*],
[edit protocols **rip**],
[edit protocols rip group *group-name* **neighbor** *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols **rip**],
[edit routing-instances *routing-instance-name* protocols rip group *group-name* **neighbor** *neighbor-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the type of authentication for RIP route queries received on an interface.

Default If you do not include this statement and the **authentication-key** statement, RIP authentication is disabled.

Options *type*—Authentication type:

- **md5**—Use the MD5 algorithm to create an encoded checksum of the packet. The encoded checksum is included in the transmitted packet. The receiving routing device uses the authentication key to verify the packet, discarding it if the digest does not match. This algorithm provides a more secure authentication scheme.
- **none**—Disable authentication. If **none** is configured, the configured authentication key is ignored.
- **simple**—Use a simple password. The password is included in the transmitted packet, which makes this method of authentication relatively insecure. The password can be from 1 through 16 contiguous letters or digits long.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **authentication-key on page 1482**
- Configuring Authentication for RIP

autonomous-system

Syntax

```
autonomous-system autonomous-system <asdot-notation> <loops number> {
  independent-domain;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Specify the routing device's AS number. The numeric range provides BGP support for 4-byte AS numbers as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*.

You can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: *<16-bit high-order value in decimal>.<16-bit low-order value in decimal>*. For example, the 4-byte AS number of 65,546 in plain-number format is represented as 1.10 in the AS-dot notation format.

Options

autonomous-system—AS number. Use a number assigned to you by the Network Information Center (NIC).

Range: 1 through 4,294,967,295 ($2^{32} - 1$) in plain-number format

Range: 0.0 through 65535.65535 in AS-dot notation format

asdot-notation—(Optional) Display the configured 4-byte autonomous system number in the AS-dot notation format.

Default: Even if a 4-byte AS number is configured in the AS-dot notation format, the default is to display the AS number in the plain-number format.

number—(Optional) Maximum number of times this AS number can appear in an AS path.

Range: 1 through 10

Default: 1 (AS number can appear once)

NOTE: When you specify the same AS number in more than one routing instance on the local routing device, you must configure the same number of loops for the AS number in each instance. For example, if you configure a value of 3 for the loops statement in a VRF routing instance that uses the same AS number as that of the master instance, you must also configure a value of 3 loops for the AS number in the master instance.

Use the independent-domain option if the loops statement must be enabled only on a subset of routing instances.

The remaining statement is explained separately.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- independent-domain
- Configuring AS Numbers for BGP

backup-pe-group

Syntax

```
backup-pe-group group-name {
    backups [ addresses ];
    local-address address;
}
```

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information
Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Configure a backup provider edge (PE) group for ingress PE redundancy when point-to-multipoint label-switched paths (LSPs) are used for multicast distribution.

Options
group-name—Name of the group for PE backups.

The remaining statements are explained separately.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Ingress PE Redundancy

backups

Syntax backups [*addresses*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast **backup-pe-group** *group-name*],
[edit logical-systems *logical-system-name* routing-options multicast **backup-pe-group** *group-name*],
[edit routing-instances *routing-instance-name* routing-options multicast **backup-pe-group** *group-name*],
[edit routing-options multicast **backup-pe-group** *group-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the address of backup PEs for ingress PE redundancy when point-to-multipoint label-switched paths (LSPs) are used for multicast distribution.

Options ***addresses***—Addresses of other PEs in the backup group.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Ingress PE Redundancy

bandwidth

Syntax bandwidth (*bps* | adaptive);

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast **flow-map**],
[edit logical-systems *logical-system-name* routing-options multicast **flow-map**],
[edit routing-instances *routing-instance-name* routing-options multicast **flow-map**],
[edit routing-options multicast **flow-map**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the bandwidth property for multicast flow maps.

Options **adaptive**—Specify that the bandwidth is measured for the flows that are matched by the flow map.

bps—Bandwidth, in bits per second, for the flow map.

Range: 0 through any amount of bandwidth

Default: 2 Mbps

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Defining Bandwidth for Multicast Flows

bandwidth-based-metrics

Syntax

```
bandwidth-based-metrics {
    bandwidth value;
    metric number;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols (ospf | ospf3) area *area-id* interface *interface-name*],
[edit logical-systems *logical-system-name* protocols ospf area *area-id* interface *interface-name* topology *topology-name*],
[edit logical-systems *logical-system-name* protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* interface *interface-name*],
[edit logical-systems *logical-system-name* protocols ospf area *area-id* interface *interface-name* topology *topology-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instances* protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* interface *interface-name*],
[edit protocols (ospf | ospf3) area *area-id* interface *interface-name*],
[edit protocols ospf area *area-id* interface *interface-name* topology *topology-name*],
[edit protocols ospf3 realm (ivp4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols ospf area *area-id* interface *interface-name* topology *topology-name*],
[edit routing-instances *routing-instance-name* protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* interface *interface-name*]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Specify a set of bandwidth threshold values and associated metric values for an OSPF interface or for a topology on an OSPF interface. When the bandwidth of an interface changes, the Junos OS automatically sets the interface metric to the value associated with the appropriate bandwidth threshold value.

Options

bandwidth *value*—Specify the bandwidth threshold in bits per second.

Range: 9600 through 1,000,000,000,000,000

metric *number*—Specify a metric value to associate with a specific bandwidth value.

Range: 1 through 65,535

NOTE: You must also configure a static metric value for the OSPF interface or topology with the metric statement. The Junos OS uses this value to calculate the cost of a route from the OSPF interface or topology if the bandwidth for the interface is higher than of any bandwidth threshold values configured for bandwidth-based metrics.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **metric on page 1611**
- Dynamically Adjusting OSPF Interface Metrics Based on Bandwidth

bfd-liveness-detection

Syntax

```
bfd-liveness-detection {
  authentication {
    algorithm algorithm-name;
    key-chain key-chain-name;
    <loose-check>;
  }
  detection-time {
    threshold milliseconds;
  }
  holddown-interval milliseconds;
  minimum-interval milliseconds;
  minimum-receive-interval milliseconds;
  multiplier number;
  no-adaptation;
  transmit-interval {
    threshold milliseconds;
    minimum-interval milliseconds;
  }
  version (1 | automatic);
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp group *group-name*],
[edit logical-systems *logical-system-name* protocols bgp group *group-name* neighbor *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address*],
[edit protocols bgp],
[edit protocols bgp group *group-name*],
[edit protocols bgp group *group-name* neighbor *address*],
[edit routing-instances *routing-instance-name* protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address*]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure bidirectional failure detection timers and authentication.

For IBGP and multihop EBGP support, configure the **bfd-liveness-detection** statement at the global **[edit bgp protocols]** hierarchy level. You can also configure IBGP and multihop support for a routing instance or a logical system.

Options

authentication algorithm *algorithm-name* —Configure the algorithm used to authenticate the specified BFD session: **simple-password**, **keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**.

authentication key-chain ***key-chain-name***—Associate a security key with the specified BFD session using the name of the security keychain. The keychain name must match one of the keychains configured in the **authentication-key-chains key-chain** statement at the **[edit security]** hierarchy level.

authentication loose-check—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication may not be configured at both ends of the BFD session.

detection-time threshold ***milliseconds***—Configure a threshold. When the BFD session detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

holddown-interval ***milliseconds***—Configure an interval specifying how long a BFD session must remain up before a state change notification is sent.

Range: 0 through 255,000

Default: 0

NOTE: You can configure the holddown-interval option only for EBGP peers.

minimum-interval ***milliseconds***—Configure the minimum intervals at which the local routing device transmits hello packets and then expects to receive a reply from a neighbor with which it has established a BFD session.

Range: 1 through 255,000

minimum-receive-interval ***milliseconds***— Configure only the minimum interval at which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session.

Range: 1 through 255,000

multiplier ***number***—Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

Range: 1 through 255

Default: 3

no-adaptation—Configure BFD sessions not to adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable to not to have BFD adaptation enabled in your network.

transmit-interval threshold ***milliseconds***—Configure a threshold. When the BFD session transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

transmit-interval minimum-interval ***milliseconds***—Configure only the minimum interval at which the local routing device transmits hello packets to a neighbor with which it has established a BFD session.

Range: 1 through 255,000

version—Configure the BFD version to detect.

Range: **1** or **automatic** (autodetect the BFD version)

Default: **automatic**

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring BFD for BGP
- Configuring BFD Authentication for BGP

bfd-liveness-detection

Syntax

```
bfd-liveness-detection {
  authentication {
    algorithm algorithm-name;
    key-chain key-chain-name;
    loose-check;
  }
  detection-time {
    threshold milliseconds;
  }
  minimum-interval milliseconds;
  minimum-receive-interval milliseconds;
  no-adaptation;
  transmit-interval {
    threshold milliseconds;
    minimum-interval milliseconds;
  }
  multiplier number;
  version (1 | automatic);
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols isis **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis **interface** *interface-name*],
[edit protocols isis **interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols isis **interface** *interface-name*]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure bidirectional failure detection timers and authentication.

Options

authentication algorithm *algorithm-name* —Configure the algorithm used to authenticate the specified BFD session: **simple-password**, **keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**.

authentication key-chain *key-chain-name*—Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the **authentication-key-chains key-chain** statement at the **[edit security]** hierarchy level.

authentication loose-check—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication may not be configured at both ends of the BFD session.

detection-time threshold *milliseconds*—Configure a threshold. When the BFD session detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

minimum-interval *milliseconds*—Configure the minimum intervals at which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session.

Range: 1 through 255,000

minimum-receive-interval ***milliseconds***—Configure only the minimum interval at which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session.

Range: 1 through 255,000

multiplier ***number***—Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

Range: 1 through 255

Default: 3

no-adaptation—Specify that BFD sessions not adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

transmit-interval threshold ***milliseconds***—Configure a threshold. When the BFD session transmit interval adapts to a value greater than the threshold, a single trap and a single system log message are sent. The interval threshold must be greater than the minimum transmit interval.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

transmit-interval minimum-interval ***milliseconds***—Configure only the minimum interval at which the routing device sends hello packets to a neighbor with which it has established a BFD session.

Range: 1 through 255,000

version—Specify the BFD version to detect.

Range: **1** (BFD version 1), or **automatic** (autodetection)

Default: **automatic**

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring BFD for IS-IS
- Configuring BFD Authentication for IS-IS

bfd-liveness-detection

Syntax

```
bfd-liveness-detection {
  authentication {
    algorithm algorithm-name;
    key-chain key-chain-name;
    loose-check;
  }
  detection-time {
    threshold milliseconds;
  }
  full-neighbors-only
  minimum-interval milliseconds;
  minimum-receive-interval milliseconds;
  no-adaptation;
  transmit-interval {
    threshold milliseconds;
    minimum-interval milliseconds;
  }
  multiplier number;
  version (1 | automatic);
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols (ospf | ospf3) area *area-id* **interface** *interface-name*],
[edit logical-systems *logical-system-name* protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **interface** *interface-name*],
[edit protocols (ospf | ospf3) area *area-id* **interface** *interface-name*],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* **interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **interface** *interface-name*]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure bidirectional failure detection timers and authentication.

Options

authentication algorithm *algorithm-name* —Configure the algorithm used to authenticate the specified BFD session: **simple-password**, **keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, or **meticulous-keyed-sha-1**.

authentication key-chain *key-chain-name*—Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the **authentication-key-chains key-chain** statement at the **[edit security]** hierarchy level.

authentication loose-check—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication may not be configured at both ends of the BFD session.

detection-time threshold ***milliseconds***—Configure a threshold. When the BFD session detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

full-neighbors-only—Establish BFD sessions only for OSPF neighbors in the full state. The default behavior is to establish BFD sessions for all OSPF neighbors.

minimum-interval ***milliseconds***—Configure the minimum intervals at which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session.

Range: 1 through 255,000 milliseconds

minimum-receive-interval ***milliseconds***—Configure only the minimum interval at which the routing device expects to receive a reply from a neighbor with which it has established a BFD session.

Range: 1 through 255,000 milliseconds

multiplier ***number***—Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

Range: 1 through 255

Default: 3

no-adaptation—Specify that BFD sessions should not adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

transmit-interval threshold ***milliseconds***—Configure a threshold. When the BFD session transmit interval adapts to a value greater than the threshold, a single trap and a single system log message are sent. The interval threshold must be greater than the minimum transmit interval.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

transmit-interval minimum-interval ***milliseconds***—Configure the minimum interval at which the routing device transmits hello packets to a neighbor with which it has established a BFD session.

Range: 1 through 255,000

version—Specify the BFD version to detect.

Range: **1** (BFD version 1) or **automatic** (autodetect version)

Default: **automatic**

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring BFD for OSPF
- Configuring BFD Authentication for OSPF

bfd-liveness-detection

Syntax

```
bfd-liveness-detection {
  authentication {
    algorithm algorithm-name;
    key-chain key-chain-name;
    <loose-check>;
  }
  detection-time {
    threshold milliseconds;
  }
  minimum-interval milliseconds;
  minimum-receive-interval milliseconds;
  transmit-interval {
    threshold milliseconds;
    minimum-interval milliseconds;
  }
  multiplier number;
  no-adaptation;
  version (1 | automatic);
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols rip **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip group *group-name* **neighbor** *neighbor-name*]
[edit protocols rip **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols rip group *group-name* **neighbor** *neighbor-name*]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure bidirectional failure detection timers and authentication.

Options

authentication algorithm ***algorithm-name*** —Configure the algorithm used to authenticate the specified BFD session: **simple-password**, **keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, or **meticulous-keyed-sha-1**.

authentication key-chain ***key-chain-name***—Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the **authentication-key-chains key-chain** statement at the **[edit security]** hierarchy level.

authentication loose-check—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication may not be configured at both ends of the BFD session.

detection-time threshold ***milliseconds***—Configure a threshold. When the BFD session detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

minimum-interval ***milliseconds***—Configure the minimum intervals at which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session.

Range: 1 through 255,000 milliseconds

minimum-receive-interval ***milliseconds***—Configure only the minimum interval at which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session.

Range: 1 through 255,000 milliseconds

multiplier ***number***—Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

Range: 1 through 255

Default: 3

no-adaptation—Configure BFD sessions not to adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

transmit-interval threshold ***milliseconds***—Configure a threshold. When the BFD session transmit interval adapts to a value greater than the threshold, a single trap and a single system log message are sent. The interval threshold must be greater than the minimum transmit interval.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

transmit-interval minimum-interval ***milliseconds***—Configure only a minimum interval at which the local routing device transmits hello packets to a neighbor.

Range: 1 through 255,000

version—Specify the BFD version to detect.

Range: (BFD version 1), or **automatic** (autodetect the version)

Default: **automatic**

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring BFD for RIP
- Configuring BFD Authentication for RIP

bfd-liveness-detection

Syntax

```
bfd-liveness-detection {
  authentication {
    algorithm algorithm-name;
    key-chain key-chain-name;
    loose-check;
  }
  detection-time {
    threshold milliseconds;
  }
  holddown-interval milliseconds;
  local-address ip-address;
  minimum-interval milliseconds;
  minimum-receive-interval milliseconds;
  minimum-receive-ttl number;
  multiplier number;
  neighbor address;
  no-adaptation;
  transmit-interval {
    threshold milliseconds;
    minimum-interval milliseconds;
  }
  version (1 | automatic);
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* static route *destination-prefix*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options static route *destination-prefix*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* static route *destination-prefix*],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],
[edit logical-systems *logical-system-name* routing-options static route *destination-prefix*],
[edit logical-systems *logical-system-name* routing-options static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* static route *destination-prefix*],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],
[edit routing-instances *routing-instance-name* routing-options static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],
[edit routing-instances *routing-instance-name* routing-options static route *destination-prefix*],
[edit routing-options rib *routing-table-name* static route *destination-prefix*],
[edit routing-options rib *routing-table-name* static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)],
[edit routing-options static route *destination-prefix*],

[edit routing-options static route *destination-prefix* qualified-next-hop (*interface-name* | *address*)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure bidirectional failure detection timers and authentication criteria for static routes.

Options

authentication algorithm ***algorithm-name*** —Configure the algorithm used to authenticate the specified BFD session: **simple-password**, **keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, or **meticulous-keyed-sha-1**.

authentication key-chain ***key-chain-name***—Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the **authentication-key-chains key-chain** statement at the **[edit security]** hierarchy level.

authentication loose-check—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication may not be configured at both ends of the BFD session.

detection-time threshold ***milliseconds***—Configure a threshold. When the Bidirectional Forwarding Detection (BFD) protocol session detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

holddown-interval ***milliseconds***—Configure an interval specifying how long a BFD session must remain up before a state change notification is sent.

Range: 0 through 255,000

Default: 0

local-address ***ip-address***—Enable a multihop BFD session and configure the source address for the BFD session.

minimum-interval ***milliseconds***—Configure the minimum intervals at which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session.

Range: 1 through 255,000

minimum-receive-interval ***milliseconds***—Configure the minimum interval at which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session.

Range: 1 through 255,000

minimum-receive-ttl ***number***—Configure the time-to-live (TTL) for the multihop BFD session.

Range: 1 through 255

Default: 255

multiplier ***number***—Configure number of hello packets not received by the neighbor that causes the originating interface to be declared down.

Range: 1 through 255

Default: 3

neighbor ***address***—Configure a next-hop address for the BFD session for a next hop specified as an interface name.

no-adaptation—Specify for BFD sessions not to adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

transmit-interval threshold *milliseconds*—Configure a threshold. When the BFD session transmit interval adapts to a value greater than the threshold, a single trap and a single system log message are sent. The interval threshold must be greater than the minimum transmit interval.

Range: 0 through 4,294,967,295

transmit-interval minimum-interval *milliseconds*—Configure the minimum interval at which the local routing device transmits hello packets to a neighbor with which it has established a BFD session.

Range: 1 through 255,000

version—Configure the BFD protocol version to detect.

Range: 1 or **automatic**

Default: **automatic** (autodetect the BFD protocol version)

The remaining statements are explained separately.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Bidirectional Forwarding Detection
- Configuring BFD Authentication for Static Routes

bgp

Syntax
bgp { ... }

Hierarchy Level
[edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit protocols],
[edit routing-instances *routing-instance-name* protocols]

Release Information
Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Enable BGP on the routing device or for a routing instance.

Default
BGP is disabled.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Enabling BGP

bgp-orf-cisco-mode

Syntax bgp-orf-cisco-mode;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols bgp **outbound-route-filter**],
[edit logical-systems *logical-system-name* protocols bgp group *group-name* **outbound-route-filter**],
[edit logical-systems *logical-system-name* protocols bgp group *group-name* neighbor *address* **outbound-route-filter**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **outbound-route-filter**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **outbound-route-filter**,
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address* **outbound-route-filter**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options **outbound-route-filter**],
[edit logical-systems *logical-system-name* routing-options **outbound-route-filter**],
[edit protocols bgp **outbound-route-filter**],
[edit protocols bgp group *group-name* **outbound-route-filter**],
[edit protocols bgp group *group-name* neighbor *address* **outbound-route-filter**],
[edit routing-instances *routing-instance-name* protocols bgp **outbound-route-filter**],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **outbound-route-filter**],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address* **outbound-route-filter**],
[edit routing-instances *routing-instance-name* routing-options **outbound-route-filter**],
[edit routing-options **outbound-route-filter**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable interoperability with routing devices that use the vendor-specific outbound route filter compatibility code of 130 and code type of 128.

NOTE: To enable interoperability for all BGP peers configured on the routing device, include the statement at the [edit routing-options outbound-route-filter] hierarchy level.

Default Disabled

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Applying Filters Provided by BGP Peers to Outbound Routes

bmp

Syntax

```
bmp {
  memory limit bytes;
  station-address (ip-address | name);
  station-port port-number;
  statistics-timeout seconds;
}
```

Hierarchy Level [edit routing-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the BGP Monitoring Protocol (BMP), which enables the routing device to collect data from the BGP **Adjacency-RIB-In** routing tables and periodically send that data to a monitoring station.

Options

memory-limit *bytes*—(Optional) Specify a threshold at which to stop collecting BMP data if the limit is exceeded.

Default: 10 MB

Range: 1,048,576 through 52,428,800

station-address (*ip-address* | *name*)—Specify the IP address or a valid URL for the monitoring where BMP data should be sent.

station-port *port-number*—Specify the port number of the monitoring station to use when sending BMP data.

statistics-timeout *seconds*—(Optional) Specify how often to send BMP data to the monitoring station.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring the BGP Monitoring Protocol

brief

Syntax (brief | full);

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options (**aggregate** | **generate**) (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* (**aggregate** | **generate**) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options (**aggregate** | **generate**) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* (**aggregate** | **generate**) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options (**aggregate** | **generate**) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* (**aggregate** | **generate**) (defaults | route)],
[edit routing-options (**aggregate** | **generate**) (defaults | route)],
[edit routing-options rib *routing-table-name* (**aggregate** | **generate**) (defaults | route)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure all AS numbers from all contributing paths to be included in the aggregate or generated route's path.

- **brief**—Include only the longest common leading sequences from the contributing AS paths. If this results in AS numbers being omitted from the aggregate route, the BGP **ATOMIC_ATTRIBUTE** path attribute is included with the aggregate route.
- **full**—Include all AS numbers from all contributing paths in the aggregate or generated route's path.

Default **full**

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **aggregate on page 1470**
- **generate on page 1543**
- Configuring Aggregate Routes
- Configuring Generated Routes

centralized

Syntax centralized;

Hierarchy Level [edit protocols lacp **ppm**]

Release Information Statement introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Disable distributed periodic packet management (PPM) processing for Link Aggregation Control Protocol (LACP) packets and run all PPM processing for LACP packets on the Routing Engine.

This statement disables distributed PPM processing for only LACP packets. You can disable distributed PPM processing for all packets that use PPM and run all PPM processing on the Routing Engine by configuring the **no-delegate-processing** statement in the **[edit routing-options ppm]** hierarchy.

BEST PRACTICE: We recommend that, generally, you disable distributed PPM only if Dell Support advises you to do so (see “Requesting Technical Support” on page lxxi). You should disable distributed PPM only if you have a compelling reason to disable it.

Default Distributed PPM processing is enabled for all packets that use PPM.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Distributed Periodic Packet Management on a J-EX Series Switch (CLI Procedure) on page 1451
- Configuring Aggregated Ethernet LACP (CLI Procedure) on page 926

check-zero

Syntax (check-zero | no-check-zero);

Hierarchy Level
[edit logical-systems *logical-system-name* protocols **rip**],
[edit logical-systems *logical-system-name* protocols rip group *group-name* **neighbor** *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **rip**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip group *group-name* **neighbor** *neighbor-name*],
[edit protocols **rip**],
[edit protocols rip group *group-name* **neighbor** *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols **rip**],
[edit routing-instances *routing-instance-name* protocols rip group *group-name* **neighbor** *neighbor-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Check whether the reserved fields in a RIP packet are zero:

- **check-zero**—Discard version 1 packets that have nonzero values in the reserved fields and version 2 packets that have nonzero values in the fields that must be zero. This default behavior implements the RIP version 1 and version 2 specifications.
- **no-check-zero**—Receive RIP version 1 packets with nonzero values in the reserved fields or RIP version 2 packets with nonzero values in the fields that must be zero. This is in spite of the fact that they are being sent in violation of the specifications in RFC 1058 and RFC 2453.

Default **check-zero**

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Accepting RIP Packets with Nonzero Values in Reserved Fields

checksum

Syntax checksum;

Hierarchy Level [edit logical-systems *logical-system-name* protocols isis **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis **interface** *interface-name*],
[edit protocols isis **interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols isis **interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable checksum for packets on this interface. The checksum cannot be enabled with MD5 hello authentication on the same interface.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Enabling Packet Checksum on IS-IS Interfaces

cluster

Syntax cluster *cluster-identifier*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name* neighbor *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],
[edit protocols bgp],
[edit protocols bgp **group** *group-name*],
[edit protocols bgp group *group-name* **neighbor** *address*],
[edit routing-instances *routing-instance-name* protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the cluster identifier to be used by the route reflector cluster in an internal BGP group.

Options ***cluster-identifier***—IPv6 or IPv4 address to use as the cluster identifier.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **no-client-reflect on page 1632**
- Configuring BGP Route Reflection

community

Syntax community ([*community-ids*] | no-advertise | no-export | no-export-subconfed | none);

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options (**aggregate** | **generate** | **static**) (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* (**aggregate** | **generate** | **static**) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options (**aggregate** | **generate** | **static**) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* (**aggregate** | **generate** | **static**) (defaults | route)]
[edit routing-instances *routing-instance-name* routing-options (**aggregate** | **generate** | **static**) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* (**aggregate** | **generate** | **static**) (defaults | route)],
[edit routing-options (**aggregate** | **generate** | **static**) (defaults | route)],
[edit routing-options rib *routing-table-name* (**aggregate** | **generate** | **static**) (defaults | route)],

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Associate BGP community information with a static, aggregate, or generated route.

Options ***community-ids***—One or more community identifiers. The ***community-ids*** format varies according to the type of attribute that you use.

The BGP community attribute format is ***as-number:community-value***:

- ***as-number***—AS number of the community member. It can be a value from 1 through 65,535.
- ***community-value***—Identifier of the community member. It can be a number from 0 through 65,535.

For more information about BGP community attributes, see the "Configuring the Extended Communities Attribute" section in the *Junos OS Policy Framework Configuration Guide*.

For specifying the BGP community attribute only, you also can specify ***community-ids*** as one of the following well-known community names defined in RFC 1997:

- **no-advertise**—Routes containing this community name are not advertised to other BGP peers.
- **no-export**—Routes containing this community name are not advertised outside a BGP confederation boundary.
- **no-export-subconfed**—Routes containing this community name are not advertised to external BGP peers, including peers in other members' ASs inside a BGP confederation.
- **none**—Explicitly exclude BGP community information with a static route. Include this option when configuring an individual route in the **route** portion to override a **community** option specified in the **defaults** portion.

NOTE: Extended community attributes are not supported at the [edit routing-options] hierarchy level. You must configure extended communities at the [edit policy-options] hierarchy level. For information about configuring extended communities, see the *Junos OS Policy Framework Configuration Guide*.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **aggregate on page 1470**
- **generate on page 1543**
- **static on page 1708**
- Configuring Static Routes
- Configuring Aggregate Routes
- Configuring Generated Routes

confederation

Syntax
confederation *confederation-autonomous-system* members [*autonomous-systems*];

Hierarchy Level
[edit logical-systems *logical-system-name* routing-options],
[edit routing-options]

Release Information
Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Specify the routing device's confederation AS number.

Options
autonomous-system—AS numbers of the confederation members.
Range: 1 through 65,535

confederation-autonomous-system—Confederation AS number. Use one of the numbers assigned to you by the NIC.
Range: 1 through 65,535

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring AS Confederation Members

csnp-interval

Syntax csnp-interval (*seconds* | disable);

Hierarchy Level [edit logical-systems *logical-system-name* protocols isis **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis **interface** *interface-name*],
[edit protocols isis **interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols isis **interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the interval between complete sequence number (CSN) packets on a LAN interface.

Options **disable**—Do not send CSN packets on this interface.

seconds—Number of seconds between the sending of CSN packets.

Range: 1 through 65,535 seconds

Default: 10 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring the Transmission Frequency for CSNP Packets on IS-IS Interfaces

damping

Syntax damping;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name* **neighbor** *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],
[edit protocols bgp],
[edit protocols bgp **group** *group-name*],
[edit protocols bgp group *group-name* **neighbor** *address*],
[edit routing-instances *routing-instance-name* protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable route flap damping.

Default Flap damping is disabled on the routing device.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Flap Damping for BGP Routes
- *Junos OS Policy Framework Configuration Guide*

dead-interval

Syntax dead-interval *seconds*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols ospf area *area-id* peer-interface *interface-name*],
[edit logical-systems *logical-system-name* protocols (ospf | ospf3) area *area-id* **interface** *interface-name*],
[edit logical-systems *logical-system-name* protocols (ospf | ospf3) area *area-id* **virtual-link**],
[edit logical-systems *logical-system-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* **virtual-link**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **interface** *interface-name*],
[edit protocols ospf area *area-id* peer-interface *interface-name*],
[edit protocols (ospf | ospf3) area *area-id* **interface** *interface-name*],
[edit protocols (ospf | ospf3) area *area-id* **virtual-link**],
[edit protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* **interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* **virtual-link**],
[edit routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify how long OSPF waits before declaring that a neighboring routing device is unavailable. This is an interval during which the routing device receives no hello packets from the neighbor.

Options ***seconds***—Interval to wait.

Range: 1 through 65,535 seconds

Default: 40 seconds (four times the hello interval)

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **hello-interval on page 1559**
- Configuring OSPF Timers

default-lsa

Syntax

```
default-lsa {
    default-metric metric;
    metric-type type;
    type-7;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols (ospf | ospf3) area *area-id* **nssa**],
[edit logical-systems *logical-system-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **nssa**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* **nssa**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **nssa**],
[edit protocols (ospf | ospf3) area *area-id* **nssa**],
[edit protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **nssa**],
[edit routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* **nssa**],
[edit routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **nssa**]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

On area border routers only, for an NSSA, inject a default LSA with a specified metric value into the area. The default route matches any destination that is not explicitly reachable from within the area.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **nssa** on page 1642
- **stub** on page 1710
- Configuring OSPF Areas

default-metric

Syntax default-metric *metric*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols (ospf | ospf3) **area** *area-id* nssa **default-lsa**],
[edit logical-systems *logical-system-name* protocols (ospf | ospf3) **area** *area-id* **stub**],
[edit logical-systems *logical-system-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* nssa **default-lsa**],
[edit logical-systems *logical-system-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **stub**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf | ospf3) **area** *area-id* nssa **default-lsa**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf | ospf3) **area** *area-id* **stub**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* nssa **default-lsa**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **stub**],
[edit protocols (ospf | ospf3) **area** *area-id* nssa **default-lsa**],
[edit protocols (ospf | ospf3) **area** *area-id* **stub**],
[edit protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* nssa **default-lsa**],
[edit protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **stub**],
[edit routing-instances *routing-instance-name* protocols (ospf | ospf3) **area** *area-id* nssa **default-lsa**],
[edit routing-instances *routing-instance-name* protocols (ospf | ospf3) **area** *area-id* **stub**],
[edit routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* nssa **default-lsa**],
[edit routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **stub**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description On area border routers only, for a stub area, inject a default route with a specified metric value into the area. The default route matches any destination that is not explicitly reachable from within the area.

Options ***metric***—Metric value.

Range: 1 through 16,777,215

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **nssa on page 1642**
- **stub on page 1710**
- Configuring OSPF Areas

description

Syntax description *text-description*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name* **neighbor** *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],
[edit protocols bgp],
[edit protocols bgp **group** *group-name*],
[edit protocols bgp group *group-name* **neighbor** *address*],
[edit routing-instances *routing-instance-name* protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Text description of the global, group, or neighbor configuration.

Options ***text-description***—Text description of the configuration. It is limited to 126 characters.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Enabling BGP
- Configuring BGP Groups and Peers
- Configuring BGP Groups and Peers

disable

Syntax disable;

Hierarchy Level [edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable BGP on the system.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Enabling BGP

disable (IS-IS)

Syntax disable;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols **isis**],
[edit logical-systems *logical-system-name* protocols isis **interface** *interface-name*],
[edit logical-systems *logical-system-name* protocols isis interface *interface-name*level *level-number*],
[edit logical-systems *logical-system-name* protocols isis traffic-engineering],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **isis**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis interface *interface-name* level *level-number*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis traffic-engineering],
[edit protocols **isis**],
[edit protocols isis **interface** *interface-name*],
[edit protocols isis interface *interface-name* level *level-number*],
[edit protocols isis traffic-engineering],
[edit routing-instances *routing-instance-name* protocols **isis**],
[edit routing-instances *routing-instance-name* protocols isis **interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols isis interface *interface-name* level *level-number*],
[edit routing-instances *routing-instance-name* protocols isis traffic-engineering]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable IS-IS on the routing device, on an interface, or on a level. At the **[edit protocols isis traffic-engineering]** hierarchy level, disable IS-IS support for traffic engineering.

Enabling IS-IS on an interface (by including the **interface** statement at the **[edit protocols isis]** or the **[edit routing-instances *routing-instance-name* protocols isis]** hierarchy level), disabling it (by including the **disable** statement), and not actually having IS-IS run on an interface (by including the **passive** statement) are mutually exclusive states.

Default IS-IS is enabled for Level 1 and Level 2 routers on all interfaces on which an International Organization for Standardization (ISO) protocol family is enabled.

IS-IS support for traffic engineering is enabled.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- IS-IS Overview
- Configuring IS-IS Traffic Engineering Attributes
- Disabling IS-IS

disable (OSPF)

Syntax disable;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols (**ospf** | **ospf3**)],
[edit logical-systems *logical-system-name* protocols (ospf | ospf3) **area** *area-id* **interface** *interface-name*],
[edit logical-systems *logical-system-name* protocols (ospf | ospf3) **virtual-link**],
[edit logical-systems *logical-system-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit logical-systems *logical-system-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf | ospf3) **area** *area-id* **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf | ospf3) **virtual-link**],
[edit logical-systems *logical-system-name* routing-instances *routing-instances* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **interface** *interface-name*],
[edit protocols (**ospf** | **ospf3**)],
[edit protocols (ospf | ospf3) **area** *area-id* **interface** *interface-name*],
[edit protocols (ospf | ospf3) **virtual-link**],
[edit protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)],
[edit routing-instances *routing-instance-name* protocols (ospf | ospf3) **area** *area-id* **interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols (ospf | ospf3) **virtual-link**],
[edit routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable OSPF, an OSPF interface, or an OSPF virtual link.

Default The configured object is enabled (operational) unless explicitly disabled.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Minimum OSPF Configuration

disable

Syntax disable;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* **routing-options** graceful-restart],
[edit logical-systems *logical-system-name* **routing-options** graceful-restart],
[edit routing-instances *routing-instance-name* **routing-options** graceful-restart],
[edit **routing-options** graceful-restart]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable graceful restart.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Graceful Restart

discard

Syntax discard;

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options (**aggregate** | **generate**) (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* (**aggregate** | **generate**) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options (**aggregate** | **generate**) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* (**aggregate** | **generate**) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options (**aggregate** | **generate**) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* (**aggregate** | **generate**) (defaults | route)],
[edit routing-options (**aggregate** | **generate**) (defaults | route)],
[edit routing-options rib *routing-table-name* (**aggregate** | **generate**) (defaults | route)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Do not forward packets addressed to this destination. Instead, drop the packets, do not send ICMP unreachable messages to the packets' originators, and install a reject route for this destination into the routing table.

Default When an aggregate route becomes active, it is installed in the routing table with a reject next hop, which means that ICMP unreachable messages are sent.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **aggregate on page 1470**
- **generate on page 1543**
- Configuring Aggregate Routes
- Configuring Generated Routes

domain-id

Syntax domain-id *domain-id*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)],
[edit routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify a domain ID for a route. The domain ID identifies the OSPF domain from which the route originated.

Options ***domain-id***—You can specify either an IP address or an IP address and a local identifier using the following format: ***ip-address*:*local-identifier***. If you do not specify a local identifier with the IP address, the identifier is assumed to have a value of 0.

Default: If the router ID is not configured in the routing instance, the router ID is derived from an interface address belonging to the routing instance.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring OSPF Domain IDs for VPNs

domain-vpn-tag

Syntax domain-vpn-tag *number*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)],
[edit routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set a virtual private network (VPN) tag for OSPFv2 external routes generated by the provider edge (PE) router.

Options ***number***—VPN tag.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring OSPF Domain IDs for VPNs

explicit-null

Syntax explicit-null;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols mpls],
[edit logical-systems *logical-system-name* protocols bgp **family** inet labeled-unicast],
[edit logical-systems *logical-system-name* protocols bgp **family** inet6 labeled-unicast],
[edit logical-systems *logical-system-name* protocols bgp group *group-name* **family** inet labeled-unicast],
[edit logical-systems *logical-system-name* protocols bgp group *group-name* **family** inet6 labeled-unicast],
[edit logical-systems *logical-system-name* protocols bgp group *group-name* neighbor *address* **family** inet labeled-unicast],
[edit logical-systems *logical-system-name* protocols bgp group *group-name* neighbor *address* **family** inet6 labeled-unicast],
[edit logical-systems *logical-system-name* protocols ldp],
[edit logical-systems *logical-system-name* routing-instances *instance-name* protocols bgp **family** inet labeled-unicast],
[edit logical-systems *logical-system-name* routing-instances *instance-name* protocols bgp **family** inet6 labeled-unicast],
[edit logical-systems *logical-system-name* routing-instances *instance-name* protocols bgp group *group-name* **family** inet labeled-unicast],
[edit logical-systems *logical-system-name* routing-instances *instance-name* protocols bgp group *group-name* **family** inet6 labeled-unicast],
[edit logical-systems *logical-system-name* routing-instances *instance-name* protocols bgp group *group-name* neighbor *address* **family** inet labeled-unicast],
[edit logical-systems *logical-system-name* routing-instances *instance-name* protocols bgp group *group-name* neighbor *address* family inet6 labeled-unicast],
[edit logical-systems *logical-system-name* routing-instances *instance-name* protocols ldp],
[edit protocols mpls],
[edit protocols bgp **family** inet labeled-unicast],
[edit protocols bgp **family** inet6 labeled-unicast],
[edit protocols bgp group *group-name* **family** inet labeled-unicast],
[edit protocols bgp group *group-name* **family** inet6 labeled-unicast],
[edit protocols bgp group *group-name* neighbor *address* **family** inet labeled-unicast]
[edit protocols bgp group *group-name* neighbor *address* family inet6 labeled-unicast],
[edit protocols ldp],
[edit routing-instances *instance-name* protocols bgp **family** inet labeled-unicast],
[edit routing-instances *instance-name* protocols bgp **family** inet6 labeled-unicast],
[edit routing-instances *instance-name* protocols bgp group *group-name* family inet labeled-unicast],
[edit routing-instances *instance-name* protocols bgp group *group-name* **family** inet6 labeled-unicast],
[edit routing-instances *instance-name* protocols bgp group *group-name* neighbor *address* **family** inet labeled-unicast],
[edit routing-instances *instance-name* protocols bgp group *group-name* neighbor *address* **family** inet6 labeled-unicast],
[edit routing-instances *instance-name* protocols ldp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Advertise label 0 to the egress routing device of an LSP.

Default If you do not include the **explicit-null** statement in the configuration, label 3 (implicit null) is advertised.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Advertising Explicit Null Labels to BGP Peers

export

Syntax export [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name* **neighbor** *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],
[edit protocols bgp],
[edit protocols bgp **group** *group-name*],
[edit protocols bgp group *group-name* **neighbor** *address*],
[edit routing-instances *routing-instance-name* protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply one or more policies to routes being exported from the routing table into BGP.

Options ***policy-names***—Name of one or more policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **import on page 1567**
- Applying Policies to BGP Routes
- *Junos OS Policy Framework Configuration Guide*

export

Syntax export [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols isis],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis],
[edit protocols isis],
[edit routing-instances *routing-instance-name* protocols isis]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply one or more policies to routes being exported from the routing table into IS-IS.

Options ***policy-names***—Name of one or more policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Applying Policies to Routes Exported to IS-IS
- *Junos OS Policy Framework Configuration Guide*
- *Junos OS Interfaces and Routing Configuration Guide*

export

Syntax

export [*policy--names*];

Hierarchy Level

[edit logical-systems *logical-system-name* protocols (**ospf** | **ospf3**)],
[edit logical-systems *logical-system-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit protocols (**ospf** | **ospf3**)],
[edit protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)],
[edit routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Apply one or more policies to routes being exported from the routing table into OSPF.

Options

policy-names—Name of one or more policies.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Applying Policies to OSPF Routes
- *Junos OS Policy Framework Configuration Guide*

export

Syntax export [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols rip **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip **group** *group-name*],
[edit protocols rip **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols rip **group** *group-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply a policy to routes being exported to the neighbors.

Options ***policy-names***—Name of one or more policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **import on page 1569**
- Configuring Group-Specific RIP Properties
- *Junos OS Policy Framework Configuration Guide*

export

Syntax export [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols ripng **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng group *group-name*],
[edit protocols ripng **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols ripng group *group-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply a policy or list of policies to routes being exported to the neighbors.

Options ***policy-names***—Name of one or more policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **import on page 1570**
- Configuring Group-Specific RIPng Properties

export

Syntax export [*policy--names*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options **forwarding-table**],
[edit logical-systems *logical-system-name* routing-options **forwarding-table**],
[edit routing-instances *routing-instance-name* routing-options **forwarding-table**],
[edit routing-options **forwarding-table**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply one or more policies to routes being exported from the routing table into the forwarding table.

Options ***policy-name***—Name of one or more policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Per-Packet Load Balancing
- *Junos OS Policy Framework Configuration Guide*

export-rib

Syntax export-rib *routing-table-name*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options **rib-group** *group-name*],
[edit logical-systems *logical-system-name* routing-options passive *group-name*],
[edit routing-instances *routing-instance-name* routing-options **rib-group** *group-name*],
[edit routing-options passive *group-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Name of the routing table from which the Junos OS should export routing information.

Options ***routing-table-name***—Routing table group name.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **import-rib on page 1572**
- passive
- Creating Routing Table Groups

external-preference

Syntax external-preference *preference*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols isis **level** *level-number*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis **level** *level-number*],
[edit protocols isis **level** *level-number*],
[edit routing-instances *routing-instance-name* protocols isis **level** *level-number*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the preference of external routes.

Options *preference*—Preference value.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

Default: 15 (for Level 1 internal routes), 18 (for Level 2 internal routes), 160 (for Level 1 external routes), 165 (for Level 2 external routes)

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **preference on page 1661**
- Configuring Preference Values for IS-IS Routes

external-preference

Syntax external-preference *preference*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols (**ospf** | **ospf3**)],
[edit logical-systems *logical-system-name* protocols ospf3 **realm** (ip4-unicast | ipv4-multicast | ipv6-multicast)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit protocols (**ospf** | **ospf3**)],
[edit protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)],
[edit routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast } ipv4-multicast | ipv6-multicast)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set the route preference for OSPF external routes.

Options *preference*—Preference value.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

Default: 150

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **preference on page 1662**
- Configuring Preference Values for OSPF Routes

family

Syntax

```
family {
  (inet | inet6 | inet-vpn | inet6-vpn | iso-vpn) {
    (any | flow | labeled-unicast | multicast | unicast) {
      accepted-prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
      }
      <loops number>;
      prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
      }
      rib-group group-name;
    }
    flow {
      no-validate policy-name;
    }
    labeled-unicast {
      accepted-prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
      }
      aggregate-label {
        community community-name:
      }
      explicit-null {
        connected-only;
      }
      prefix-limit {
        maximum number;
        teardown <percentage> <idle-timeout (forever | minutes)>;
      }
      resolve-vpn;
      rib inet.3;
      rib-group group-name;
    }
  }
  route-target {
    accepted-prefix-limit {
      maximum number;
      teardown <percentage> <idle-timeout (forever | minutes)>;
    }
    advertise-default;
    external-paths number;
    prefix-limit {
      maximum number;
      teardown <percentage> <idle-timeout (forever | minutes)>;
    }
  }
  (inet-mdt | inet-mvpn | inet6-mvpn | l2-vpn) {
    signaling {
      accepted-prefix-limit {
```

```
            maximum number;
            teardown <percentage> <idle-timeout (forever | minutes)>;
          }
          <loops number>;
          prefix-limit {
            maximum number;
            teardown <percentage> <idle-timeout (forever | minutes)>;
            }
          rib-group group-name
        }
      }
    }
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name* **neighbor** *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],
[edit protocols bgp],
[edit protocols bgp **group** *group-name*],
[edit protocols bgp group *group-name* **neighbor** *address*],
[edit routing-instances *routing-instance-name* protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Enable multiprotocol BGP (MP-BGP) by configuring BGP to carry network layer reachability information (NLRI) for address families other than unicast IPv4, to specify MP-BGP to carry NLRI for the IPv6 address family, or to carry NLRI for VPNs.

Options

any—Configure the family type to be both unicast and multicast.

inet—Configure NLRI parameters for IPv4.

inet6—Configure NLRI parameters for IPv6.

inet-mdt—Configure NLRI parameters for the multicast distribution tree (MDT) subaddress family identifier (SAFI) for IPv4 traffic in Layer 3 VPNs.

inet-mvpn—Configure NLRI parameters for IPv4 for multicast VPNs.

inet6-mvpn—Configure NLRI parameters for IPv6 for multicast VPNs.

inet-vpn—Configure NLRI parameters for IPv4 for Layer 3 VPNs.

inet6-vpn—Configure NLRI parameters for IPv6 for Layer 3 VPNs.

iso-vpn—Configure NLRI parameters for IS-IS for Layer 3 VPNs.

l2-vpn—Configure NLRI parameters for IPv4 for MPLS-based Layer 2 VPNs and VPLS.

labeled-unicast—Configure the family type to be labeled-unicast. This means that the BGP peers are being used only to carry the unicast routes that are being used by labeled-unicast for resolving the labeled-unicast routes. This statement is supported only with **inet** and **inet6**.

loops ***number***—(Optional) Specify the maximum number of times that the AS number can appear in the AS path received from a BGP peer for the specified address family. For ***number***, include a value from 1 through 10.

NOTE: When you configure the **loops** statement for a specific BGP address family, that value is used to evaluate the AS path for routes received by a BGP peer for the specified address family rather than the loops value configured for the global AS number.

multicast—Configure the family type to be multicast. This means that the BGP peers are being used only to carry the unicast routes that are being used by multicast for resolving the multicast routes.

unicast—Configure the family type to be unicast. This means that the BGP peers only carry the unicast routes that are being used for unicast forwarding purposes.

Default: unicast

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- autonomous-system on page 1487
- local-as on page 1597
- Enabling Multiprotocol BGP

fate-sharing

Syntax

```
fate-sharing {
  group group-name {
    cost value;
    from address <to address>;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit routing-options],
[edit routing-instances *routing-instance-name* routing-options]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Specify groups of objects that share characteristics resulting in backup paths to be used if primary paths fail. All objects are treated as **/32** host addresses. You specify one or more objects within a group. The objects can be LAN interfaces, router IDs, or point-to-point links. The sequence is insignificant.

Options

cost *value*—Cost assigned to the group.

Range: 1 through 65,535

Default: 1

from *address*—Address of the router or address of the LAN/NBMA interface. For example, an Ethernet network with four hosts in the same fate-sharing group would require you to list all four of the separate **from** addresses in the group.

group *group-name*—Each fate-sharing group must have a name, which can have a maximum of 32 characters, including letters, numbers, periods (.), and hyphens (-). You can define up to 512 groups.

to *address*—(Optional) Address of egress router. For point-to-point link objects, you must specify both a **from** and a **to** address.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring the Ingress Router for MPLS-Signaled LSPs

flow

Syntax

```
flow {
  route name {
    match {
      match-conditions;
    }
    term-order (legacy | standard);
    then {
      actions;
    }
  }
  validation {
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>;
      flag flag <flag-modifier> <disable>;
    }
  }
}
```

Hierarchy Level [edit routing-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a flow route.

Default legacy

Options

actions—An action to take if conditions match.

match-conditions—Match packets to these conditions.

route *name*—Name of the flow route.

standard—Specify to use version 7 or later of the flow-specification algorithm.

term-order (legacy | standard)—Specify the version of the flow-specification algorithm.

- **legacy**—Use version 6 of the flow-specification algorithm.
- **standard**—Use version 7 of the flow-specification algorithm.

then—Actions to take on matching packets.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Flow Routes

flow-map

Syntax

```
flow-map flow-map-name {
  bandwidth (bps | adaptive);
  forwarding-cache {
    timeout (never | minutes);
  }
  policy [ policy-names ];
  redundant-sources [ addresses ];
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure multicast flow maps.

Options ***flow-map-name***—Name of the flow-map.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Creating a Multicast Flow Map

forwarding-cache (Flow Maps)

Syntax

```
forwarding-cache {
  timeout (minutes | never );
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast **flow-map** *flow-map-name*],
[edit logical-systems *logical-system-name* routing-options multicast **flow-map** *flow-map-name*],
[edit routing-instances *routing-instance-name* routing-options multicast **flow-map** *flow-map-name*],
[edit routing-options multicast **flow-map** *flow-map-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure multicast forwarding cache properties for the flow map.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

forwarding-cache (Multicast)

Syntax

```
forwarding-cache {
    threshold suppress value <reuse value>;
    timeout minutes;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure multicast forwarding cache properties. These properties include threshold suppression and reuse limits and timeout values.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring General Multicast Forwarding Cache Properties

forwarding-table

Syntax

```
forwarding-table {
    export [ policy--names ];
    (indirect-next-hop | no-indirect-next-hop);
    unicast-reverse-path (active-paths | feasible-paths);
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-options],
[edit routing-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure information about the routing device's forwarding table.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Per-Packet Load Balancing

generate

Syntax

```
generate {
  defaults {
    generate-options;
  }
  route destination-prefix {
    policy policy-name;
    generate-options;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options **rib** *routing-table-name*],
[edit routing-options],
[edit routing-options **rib** *routing-table-name*]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure generated routes, which are used as routes of last resort.

Options

generate-options—Additional information about generated routes, which is included with the route when it is installed in the routing table. Specify zero or more of the following options in ***generate-options***. Each option is explained separately.

- **(active | passive);**
- **as-path <***as-path***> <origin (egp | igp | incomplete)> <atomic-aggregate> <aggregator** *as-number in-address***>;**
- **community [** *community-ids* **];**
- **discard;**
- **(brief | full);**
- **(metric | metric2 | metric3 | metric4)** *value* **<type** *type***>;**
- **(preference | preference2 | color | color2)** *preference* **<type** *type***>;**
- **tag** *string***;**

defaults—Specify global generated route options. These options only set default attributes inherited by all newly created generated routes. These are treated as global defaults and apply to all the generated routes you configure in the **generate** statement. This part of the **generate** statement is optional.

route ***destination-prefix***—Configure a non-default generated route:

- **default**—For the default route to the destination. This is equivalent to specifying an IP address of **0.0.0.0/0**.

- ***destination-prefix/prefix-length***—/***destination-prefix*** is the network portion of the IP address, and ***prefix-length*** is the destination prefix length.

The **policy** statement is explained separately.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Generated Routes

graceful-restart

Syntax

```
graceful-restart {
  disable;
  restart-timeseconds;
  stale-routes-time seconds;
}
```

Hierarchy Level
[edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name* **neighbor** *address*],
[edit protocols bgp],
[edit protocols bgp **group** *group-name*],
[edit protocols bgp group *group-name* **neighbor** *address*]

Release Information
Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Configure graceful restart for BGP.

Options
disable—Disable graceful restart for BGP.

restart-time *seconds*—Time period when the restart is expected to be complete.
Range: 1 through 600 seconds

stale-routes-time *seconds*—Maximum time that stale routes are kept during restart.
Range: 1 through 600 seconds

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Graceful Restart
- Configuring Graceful Restart for BGP
- *Junos OS High Availability Configuration Guide*

graceful-restart

Syntax

```
graceful-restart {
  disable;
  helper-disable;
  restart-duration seconds;
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols isis],
[edit protocols isis]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure graceful restart for IS-IS.

Options

disable—Disable graceful restart.

helper-disable—Disable graceful restart helper capability. Helper mode is enabled by default.

restart-duration *seconds*—Configure the time period for the restart to last, in seconds.

Range: 30 through 300 seconds

Default: 30 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Graceful Restart
- Configuring Graceful Restart for IS-IS

graceful-restart

Syntax

```
graceful-restart {
  disable;
  helper-disable;
  notify-duration seconds;
  restart-duration seconds;
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols (**ospf** | **ospf3**)],
[edit protocols (**ospf** | **ospf3**)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure graceful restart for OSPF.

Options

disable—Disable graceful restart for OSPF.

helper-disable—Disable graceful restart helper capability. Helper mode is enabled by default.

notify-duration *seconds*—Estimated time to send out purged grace LSAs over all the interfaces.

Range: 1 through 3600 seconds

Default: 30 seconds

restart-duration *seconds*—Estimated time to reacquire a full OSPF neighbor from each area.

Range: 1 through 3600 seconds

Default: 180 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Graceful Restart for OSPF and OSPFv3
- *Junos OS High Availability Configuration Guide*

graceful-restart

Syntax

```
graceful-restart {
  disable;
  restart-time seconds;
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols **rip**],
[edit protocols **rip**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure graceful restart for RIP.

Options **disable**—Disables graceful restart for RIP.

seconds—Estimated time for the restart to finish, in seconds.

Range: 1 through 600 seconds

Default: 60 seconds

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Graceful Restart
- Configuring Graceful Restart for RIP

graceful-restart

Syntax

```
graceful-restart {
  disable;
  restart-time seconds;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols ripng],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng],
[edit protocols ripng],
[edit routing-instances *routing-instance-name* protocols ripng]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure graceful restart for RIPng.

Options

disable—Disables graceful restart for RIPng.

seconds—Estimated time period for the restart to finish.

Range: 1 through 600 seconds

Default: 60 seconds

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Graceful Restart
- Configuring Graceful Restart for RIPng

graceful-restart

Syntax

```
graceful-restart {
  disable;
  restart-duration seconds;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* **routing-options**],
[edit logical-systems *logical-system-name* **routing-options**],
[edit routing-instances *routing-instance-name* **routing-options**],
[edit **routing-options**]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure graceful restart.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Graceful Restart
- *Junos OS High Availability Configuration Guide*

group

```
Syntax    group group-name {
            advertise-inactive;
            allow [ network/mask-length ];
            authentication-key key;
            cluster cluster-identifier;
            damping;
            description text-description;
            export [ policy-names ];
            family {
              (inet | inet6 | inet-vpn | inet6-vpn | l2-vpn) {
                (any | multicast | unicast | signaling) {
                  accepted-prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                  }
                  prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                  }
                  rib-group group-name;
                }
                flow {
                  no-validate policy-name;
                }
                labeled-unicast {
                  accepted-prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                  }
                  explicit-null {
                    connected-only;
                  }
                  prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                  }
                  resolve-vpn;
                  rib inet.3;
                  rib-group group-name;
                }
              }
              route-target {
                accepted-prefix-limit {
                  maximum number;
                  teardown <percentage> <idle-timeout (forever | minutes)>;
                }
                advertise-default;
                external-paths number;
                prefix-limit {
                  maximum number;
                  teardown <percentage> <idle-timeout (forever | minutes)>;
                }
```

```
    }
  }
  hold-time seconds;
  import [ policy-names ];
  ipsec-sa ipsec-sa;
  keep (all | none);
  local-address address;
  local-as autonomous-system <private>;
  local-preference local-preference;
  log-updown;
  metric-out metric;
  multihop <ttl-value>;
  multipath {
    multiple-as;
  }
  no-aggregator-id;
  no-client-reflect;
  out-delay seconds;
  passive;
  peer-as autonomous-system;
  preference preference;
  remove-private;
  tcp-mss segment-size;
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
  }
  type type;
  neighbor address {
    ... peer-specific-options ...
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Define a BGP peer group. BGP peer groups share a common type, peer autonomous system (AS) number, and cluster ID, if present. To configure multiple BGP groups, include multiple **group** statements.

By default, the group's options are identical to the global BGP options. To override the global options, include group-specific options within the **group** statement.

The **group** statement is one of the statements you must include in the configuration to run BGP on the routing device. See Minimum BGP Configuration.

Options

group-name—Name of the BGP group.

The remaining statements are explained separately.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring BGP Groups and Peers

group

Syntax

```
group group-name {
  bfd-liveness-detection {
    authentication {
      algorithm algorithm-name;
      key-chain key-chain-name;
      loose-check;
    }
    detection-time {
      threshold milliseconds;
    }
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    transmit-interval {
      threshold milliseconds;
      minimum-interval milliseconds;
    }
    multiplier number;
    version (0 | 1 | automatic);
  }
  preference number;
  metric-out metric;
  export policy;
  route-timeout seconds;
  update-interval seconds;
  neighbor neighbor-name {
    authentication-key password;
    authentication-type type;
    bfd-liveness-detection {
      authentication {
        algorithm algorithm-name;
        key-chain key-chain-name;
        loose-check;
      }
      detection-time {
        threshold milliseconds;
      }
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      transmit-interval {
        threshold milliseconds;
        minimum-interval milliseconds;
      }
      multiplier number;
      version (0 | 1 | automatic);
    }
    (check-zero | no-check-zero);
    import policy-name;
    message-size number;
    metric-in metric;
    metric-out metric;
    receive receive-options;
    route-timeout seconds;
```

```
        send send-options;
        update-interval seconds;
    }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols **rip**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **rip**],
[edit protocols **rip**],
[edit routing-instances *routing-instance-name* protocols **rip**]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure a set of RIP neighbors that share an export policy and metric. The export policy and metric govern what routes to advertise to neighbors in a given group.

Options

group-name—Name of a group, up to 16 characters long.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Group-Specific RIP Properties

group

Syntax

```
group group-name {
  export [ policy-names ];
  metric-out metric;
  preference number;
  route-timeout seconds;
  update-interval seconds;
  neighbor neighbor-name {
    import policy-name;
    metric-in metric;
    receive <none>;
    route-timeout seconds;
    send <none>;
    update-interval seconds;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols ripng],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng],
[edit protocols ripng],
[edit routing-instances *routing-instance-name* protocols ripng]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure a set of RIPng neighbors that share an export policy and metric. The export policy and metric govern what routes to advertise to neighbors in a given group.

Options

group-name—Name of a group, up to 16 characters long.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Group-Specific RIPng Properties

hello-authentication-key

Syntax hello-authentication-key *password*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols isis interface *interface-name*level *number*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis interface *interface-name* level *number*],
[edit protocols isis interface *interface-name* level *number*],
[edit routing-instances *routing-instance-name* protocols isis interface *interface-name* level *number*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure an authentication key (password) for hello packets. Neighboring routing devices use the password to verify the authenticity of packets sent from an interface. For the key to work, you also must include the **hello-authentication-type** statement.

Default By default, hello authentication is not configured on an interface. However, if IS-IS authentication is configured, the hello packets are authenticated using the IS-IS authentication type and password.

Options ***password***—Authentication password. The password can be up to 255 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (“ ”).

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **authentication-key on page 1481**
- **authentication-type on page 1485**
- **hello-authentication-type on page 1557**
- Configuring Levels on IS-IS Interfaces

hello-authentication-type

Syntax hello-authentication-type (md5 | simple);

Hierarchy Level [edit logical-systems *logical-system-name* protocols isis interface *interface-name* level *number*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis interface *interface-name* level *number*],
[edit protocols isis interface *interface-name* level *number*],
[edit routing-instances *routing-instance-name* protocols isis interface *interface-name* level *number*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable authentication on an interface for hello packets. If you enable authentication on hello packets, you must specify a password by including the **hello-authentication-key** statement.

Default By default, hello authentication is not configured on an interface. However, if IS-IS authentication is configured, the hello packets are authenticated using the IS-IS authentication type and password.

Options **md5**—Specifies Message Digest 5 as the packet verification type.

simple—Specifies simple authentication as the packet verification type.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **authentication-key on page 1481**
- **authentication-type on page 1485**
- **hello-authentication-key on page 1556**
- Configuring Levels on IS-IS Interfaces

hello-interval

Syntax hello-interval *seconds*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols isis interface *interface-name* level *level-number*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis interface *interface-name* level *level-number*],
[edit protocols isis interface *interface-name* level *level-number*],
[edit routing-instances *routing-instance-name* protocols isis interface *interface-name* level *level-number*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Frequency with which the routing device sends hello packets out of an interface, in seconds.

Options ***seconds***—Frequency of transmission for hello packets.

Range: 1 through 20,000 seconds

Default: 3 seconds (for designated intersystem [DIS] routers), 9 seconds (for non-DIS routers)

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- hold-time
- Configuring Levels on IS-IS Interfaces

hello-interval

Syntax hello-interval *seconds*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols ospf area *area-id* peer-interface *interface-name*],
[edit logical-systems *logical-system-name* protocols (ospf | ospf3) area *area-id* **interface** *interface-name*],
[edit logical-systems *logical-system-name* protocols (ospf | ospf3) area *area-id* **virtual-link**],
[edit logical-systems *logical-system-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* **virtual-link**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **interface** *interface-name*],
[edit protocols ospf area *area-id* peer-interface *interface-name*],
[edit protocols (ospf | ospf3) area *area-id* **interface** *interface-name*],
[edit protocols (ospf | ospf3) area *area-id* **virtual-link**],
[edit protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* **interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* **virtual-link**],
[edit routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify how often the routing device sends hello packets out the interface. The hello interval must be the same for all routing devices on a shared logical IP network.

Options ***seconds***—Time between hello packets, in seconds.

Range: 1 through 255 seconds

Default: 10 seconds; 120 seconds (nonbroadcast networks)

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **dead-interval on page 1518**
- Configuring OSPF Timers

hello-padding

Syntax hello-padding (adaptive | loose | strict);

Hierarchy Level [edit logical-systems *logical-system-name* protocols isis **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis **interface** *interface-name*],
[edit protocols isis **interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols isis **interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure padding on hello packets to accommodate asymmetrical maximum transfer units (MTUs) from different hosts.

Options **adaptive**—Configure padding until state of neighbor adjacency is up.

loose—Configure padding until state of adjacency is initialized.

strict—Configure padding for all adjacency states.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Enabling Padding of IS-IS Hello Packets

holddown

Syntax holddown *seconds*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols **rip**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **rip**],
[edit protocols **rip**],
[edit routing-instances *routing-instance-name* protocols **rip**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the time period the expired route is retained in the routing table before being removed.

Options ***seconds***—Estimated time to wait before making updates to the routing table.

Range: 10 through 180 seconds

Default: 180 seconds

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring RIP Timers

holddown

Syntax holddown *seconds*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols ripng],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng],
[edit protocols ripng],
[edit routing-instances *routing-instance-name* protocols ripng]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the time period the expired route is retained in the routing table before being removed.

Options ***seconds***—Estimated time to wait before making updates to the routing table.

Default: 180 seconds

Range: 10 through 180 seconds

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring RIPng Timers

hold-time

Syntax hold-time *seconds*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* vrrp-inet6-group *group-id*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the time in seconds after which a backup router with the highest priority preempts the master router.

Options ***seconds***—Hold-time period.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring VRRP for IPv6 (CLI Procedure) on page 1452

hold-time

Syntax hold-time *seconds*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name* **neighbor** *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],
[edit protocols bgp],
[edit protocols bgp **group** *group-name*],
[edit protocols bgp group *group-name* **neighbor** *address*],
[edit routing-instances *routing-instance-name* protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the hold-time value to use when negotiating a connection with the peer. The hold-time value is advertised in open packets and indicates to the peer the length of time that it should consider the sender valid. If the peer does not receive a keepalive, update, or notification message within the specified hold time, the BGP connection to the peer is closed and routing devices through that peer become unavailable.

The hold time is three times the interval at which keepalive messages are sent.

Options ***seconds***—Hold time.

Range: 20 through 65,535 seconds

Default: 90 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring the Delay Before BGP Peers Mark the Routing Device as Down

hold-time (IS-IS)

Syntax hold-time *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols isis interface *interface-name*level *level-number*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis interface *interface-name* level *level-number*],
[edit protocols isis interface *interface-name* level *level-number*],
[edit routing-instances *routing-instance-name* protocols isis interface *interface-name* level *level-number*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set the length of time a neighbor considers this router to be operative (up) after receiving a hello packet. If the neighbor does not receiver another hello packet within the specified time, it marks this routing device as inoperative (down). The hold time itself is advertised in the hello packets.

Options ***seconds***—Hold-time value, in seconds.

Range: 3 through 65,535 seconds, or 1 to send out hello packets every 333 milliseconds

Default: 9 seconds (for DIS routers), 27 seconds (for non-DIS routers; three times the default hello interval)

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **hello-interval on page 1558**
- Configuring Levels on IS-IS Interfaces

idle-after-switch-over

Syntax idle-after-switch-over (forever | *seconds*);

Hierarchy Level [edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp group *group-name*],
[edit logical-systems *logical-system-name* protocols bgp group *group-name* neighbor *address*],
[edit protocols bgp],
[edit protocols bgp group *group-name*],
[edit protocols bgp group *group-name* neighbor *address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the routing device not to automatically reestablish BGP peering sessions after a nonstop active routing (NSR) switchover. This feature is particularly useful if you are using dynamic routing policies because the dynamic database is not synchronized with the backup Routing Engine when NSR is enabled.

Options

forever—Do not reestablish a BGP peering session after an NSR switchover until the **clear bgp neighbor** command is issued.

seconds—Do not reestablish a BGP peering session after an NSR switchover until after the specified period.

Range: 1 through 4,294,967,295 ($2^{32} - 1$)

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Preventing Automatic Reestablishment of BGP Peering Sessions After NSR Switchovers
- *Junos OS Policy Framework Configuration Guide*
- *Junos OS High Availability Configuration Guide*

ignore-attached-bit

Syntax ignore-attached-bit;

Hierarchy Level [edit logical-systems *logical-system-name* protocols **isis**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **isis**],
[edit protocols **isis**],
[edit routing-instances *routing-instance-name* protocols **isis**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Ignore the attached bit on IS-IS Level 1 routers. Configuring this statement allows the routing device to ignore the attached bit on incoming Level 1 LSPs. If the attached bit is ignored, no default route, which points to the routing device which has set the attached bit, is installed.

Default The **ignore-attached-bit** statement is disabled by default.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring IS-IS

ignore-lsp-metrics

Syntax ignore-lsp-metrics;

Hierarchy Level [edit logical-systems *logical-system-name* protocols ospf **traffic-engineering shortcuts**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf **traffic-engineering shortcuts**],
[edit protocols ospf **traffic-engineering shortcuts**],
[edit routing-instances *routing-instance-name* protocols ospf **traffic-engineering shortcuts**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Ignore RSVP LSP metrics in OSPF traffic engineering shortcut calculations.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Enabling OSPF Traffic Engineering Support

import

Syntax import [*policy-names*];

Hierarchy Level
[edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name* **neighbor** *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],
[edit protocols bgp],
[edit protocols bgp **group** *group-name*],
[edit protocols bgp group *group-name* **neighbor** *address*],
[edit routing-instances *routing-instance-name* protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply one or more routing policies to routes being imported into the Junos OS routing table from BGP.

Options ***policy-names***—Name of one or more policies.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **export on page 1529**
- Applying Policies to BGP Routes
- *Junos OS Policy Framework Configuration Guide*

import

Syntax import [*policy--names*];

Hierarchy Level
[edit logical-systems *logical-system-name* protocols (**ospf** | **ospf3**)],
[edit logical-systems *logical-system-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit protocols (**ospf** | **ospf3**)],
[edit protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)],
[edit routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Filter OSPF routes from being added to the routing table.

Options ***policy-names***—Name of one or more policies.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Applying Policies to OSPF Routes
- *Junos OS Policy Framework Configuration Guide*

import

Syntax import [*policy-names*];

Hierarchy Level
[edit logical-systems *logical-system-name* protocols **rip**],
[edit logical-systems *logical-system-name* protocols rip group *group-name* **neighbor** *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **rip**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip group *group-name* **neighbor** *neighbor-name*],
[edit protocols **rip**],
[edit protocols rip group *group-name* **neighbor** *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols **rip**],
[edit routing-instances *routing-instance-name* protocols rip group *group-name* **neighbor** *neighbor-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply one or more policies to routes being imported by the local router from its neighbors.

Options ***policy-names***—Name of one or more policies.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **export on page 1532**
- Applying Policies to RIP Routes Imported from Neighbors
- *Junos OS Policy Framework Configuration Guide*

import

Syntax import [*policy-names*];

Hierarchy Level
[edit logical-systems *logical-system-name* protocols ripng],
[edit logical-systems *logical-system-name* protocols ripng group *group-name* neighbor *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng group *group-name* neighbor *neighbor-name*],
[edit protocols ripng],
[edit protocols ripng group *group-name* neighbor *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols ripng],
[edit routing-instances *routing-instance-name* protocols ripng group *group-name* neighbor *neighbor-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply one or more policies to routes being imported into the local routing device from the neighbors.

Options ***policy-names***—Name of one or more policies.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **export on page 1532**
- Applying Policies to RIPng Routes Imported from Neighbors

import

Syntax import [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options resolution **rib**],
[edit logical-systems *logical-system-name* routing-options resolution **rib**],
[edit routing-instances *routing-instance-name* routing-options resolution **rib**],
[edit routing-options resolution **rib**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify one or more import policies to use for route resolution.

Options ***policy-names***—Name of one or more import policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Route Resolution

import-policy

Syntax import-policy [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options **rib-group** *group-name*],
[edit logical-systems *logical-system-name* routing-options passive *group-name*],
[edit routing-instances *routing-instance-name* routing-options **rib-group** *group-name*],
[edit routing-options **rib-groups** *group-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply one or more policies to routes imported into the routing table group. The **import-policy** statement complements the **import-rib** statement and cannot be used unless you first specify the routing tables to which routes are being imported.

Options ***policy-names***—Name of one or more policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **export-rib on page 1533**
- passive
- Creating Routing Table Groups

import-rib

Syntax import-rib [*routing-table--names*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options **rib-group** *group-name*],
[edit logical-systems *logical-system-name* routing-options **rib-group** *group-name*],
[edit routing-instances *routing-instance-name* routing-options **rib-group** *group-name*],
[edit routing-options **rib-group** *group-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Name of the routing table into which Junos OS should import routing information. The first routing table name you enter is the primary routing table. Any additional names you enter identify secondary routing tables. When a protocol imports routes, it imports them into the primary and any secondary routing tables. If the primary route is deleted, the secondary route also is deleted. For IPv4 import routing tables, the primary routing table must be **inet.0** or **routing-instance-name.inet.0**. For IPv6 import routing tables, the primary routing table must be **inet6.0**.

You can configure an IPv4 import routing table that includes both IPv4 and IPv6 routing tables. Including both types of routing tables permits you, for example, to populate an IPv6 routing table with IPv6 addresses that are compatible with IPv4.

Options ***routing-table-names***—Name of one or more routing tables.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **export-rib on page 1533**
- passive
- Creating Routing Table Groups

include-mp-next-hop

Syntax include-mp-next-hop;

Hierarchy Level [edit logical-systems *logical-system-name* protocols bgp],
[edit protocols bgp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable multiprotocol updates to contain next-hop reachability information.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Including Next-Hop Reachability Information in Multiprotocol Updates

indirect-next-hop

Syntax (indirect-next-hop | no-indirect-next-hop);

Hierarchy Level [edit logical-systems *logical-system-name* routing-options **forwarding-table**],
[edit routing-options **forwarding-table**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable indirectly connected next hops for route convergence.

NOTE: When virtual private LAN service (VPLS) is configured on the routing device, the indirect-next-hop statement is not supported.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Enabling Indirect Next Hops

inet6-advertise-interval

Syntax inet6-advertise-interval *milliseconds*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* vrrp-inet6-group *group-id*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the interval between Virtual Router Redundancy Protocol (VRRP) IPv6 advertisement packets.

Options ***milliseconds***—Interval, in milliseconds, between advertisement packets.

Range: 100 to 40,000 ms

Default: 1 second

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring VRRP for IPv6 (CLI Procedure) on page 1452

install

Syntax (install | no-install);

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options **static** (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* **static** (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* **static** (defaults | route)],
[edit logical-systems *logical-system-name* routing-options **static** (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* **static** (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options **static** (defaults | route)],
[edit routing-options rib *routing-table-name* **static** (defaults | route)]
[edit routing-options **static** (defaults | route)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure whether the Junos OS installs all static routes into the forwarding table. Even if you configure a route so it is not installed in the forwarding table, the route is still eligible to be exported from the routing table to other protocols.

Options
install—Explicitly install all static routes into the forwarding table.

no-install—Do not install the route into the forwarding table, even if it is the route with the lowest preference.

Default: **install**

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **static on page 1708**
- Configuring Static Routes

instance-export

Syntax instance-export [*policy--names*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply one or more policies to routes being exported from a routing instance.

Options ***policy-names***—Name of one or more export policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Policy-Based Export for Routing Instances
- *Junos OS Policy Framework Configuration Guide*

instance-import

Syntax instance-import [*policy--names*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply one or more policies to routes being imported into a routing instance.

Options ***policy-names***—Name of one or more import policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Policy-Based Export for Routing Instances
- *Junos OS Policy Framework Configuration Guide*

inter-area-prefix-export

Syntax inter-area-prefix-export [*policy-names*];

Hierarchy Level
[edit logical-systems *logical-system-name* protocols ospf3 **area** *area-id*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 **area** *area-id*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 **realm** (ip4-unicast | ipv4-multicast | ipv6-multicast) **area** *area-id*
[edit protocols ospf3 **area** *area-id*],
[edit protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) **area** *area-id*],
[edit routing-instances *routing-instance-name* protocols ospf3 **area** *area-id*],
[edit routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-muticast | ipv6-multicast) **area** *area-id*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply an export policy for OSPFv3 to specify which interarea prefix link-state advertisements (LSAs) are flooded into an area.

Options ***policy-name***—Name of a policy configured at the **[edit policy-options policy-statement *policy-name* term *term-name*]** hierarchy level.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **inter-area-prefix-import on page 1578**
- Applying Policies to OSPF Routes
- *Junos OS Policy Framework Configuration Guide*

inter-area-prefix-import

Syntax inter-area-prefix-import [*policy-names*];

Hierarchy Level
[edit logical-systems *logical-system-name* protocols ospf3 **area** *area-id*],
[edit logical-systems *logical-system-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) **area** *area-id*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 **area** *area-id*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) **area** *area-id*],
[edit protocols ospf3 area *area-id*],
[edit protocols ospf3 **realm** (ip4-unicast | ipv4-multicast | ipv6-multicast)], **area** *area-id*],
[edit routing-instances *routing-instance-name* protocols ospf3 **area** *area-id*],
[edit routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) **area** *area-id*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply an import policy for OSPFv3 to specify which routes learned from an area are used to generate interarea prefixes into other areas.

Options ***policy-name***—Name of a policy configured at the **[edit policy-options policy-statement *policy-name* term *term-name*]** hierarchy level.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **inter-area-prefix-export on page 1577**
- Applying Policies to OSPF Routes
- *Junos OS Policy Framework Configuration Guide*

interface

Syntax

```
interface (all | interface-name) {
  disable;
  bfd-liveness-detection {
    authentication {
      algorithm algorithm-name;
      key-chain key-chain-name;
      loose-check;
    }
    detection-time {
      threshold milliseconds;
    }
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    transmit-interval {
      threshold milliseconds;
      minimum-interval milliseconds;
    }
    multiplier number;
  }
  checksum;
  csnp-interval (seconds | disable);
  hello-padding (adaptive | loose | strict);
  ldp-synchronization {
    disable;
    hold-time seconds;
  }
  lsp-interval milliseconds;
  mesh-group (value | blocked);
  no-adjacency-holddown;
  no-ipv4-multicast;
  no-ipv6-multicast;
  no-ipv6-unicast;
  no-unicast-topology;
  passive;
  point-to-point;
  level level-number {
    disable;
    hello-authentication-type authentication;
    hello-authentication-key key;
    hello-interval seconds;
    hold-time seconds;
    ipv4-multicast-metric number;
    ipv6-multicast-metric number;
    ipv6-unicast-metric number;
    metric metric;
    passive;
    priority number;
    te-metric metric;
  }
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols **isis**],

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis],
[edit protocols isis],
[edit routing-instances *routing-instance-name* protocols isis]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure interface-specific IS-IS properties. To configure more than one interface, include the **interface** statement multiple times.

Enabling IS-IS on an interface (by including the **interface** statement at the **[edit protocols isis]** or the **[edit routing-instances *routing-instance-name* protocols isis]** hierarchy level), disabling it (by including the **disable** statement), and not actually having IS-IS run on an interface (by including the **passive** statement) are mutually exclusive states.

Options **all**—Have the Junos OS create IS-IS interfaces automatically.

interface-name—Name of an interface. Specify the full interface name, including the physical and logical address components. For details about specifying interfaces, see the *Junos OS Network Interfaces Configuration Guide* and the *Junos OS Services Interfaces Configuration Guide*.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring of Interface-Specific IS-IS Properties

interface

Syntax

```
interface interface-name {
  disable;
  authentication key <key-id identifier>;
  bfd-liveness-detection {
    authentication {
      algorithm algorithm-name;
      key-chain key-chain-name;
      loose-check;
    }
    detection-time {
      threshold milliseconds;
    }
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    transmit-interval {
      threshold milliseconds;
      minimum-interval milliseconds;
    }
    multiplier number;
  }
  dead-interval seconds;
  demand-circuit;
  hello-interval seconds;
  ipsec-sa name;
  interface-type type;
  ldp-synchronization {
    disable;
    hold-time seconds;
  }
  metric metric;
  neighbor address <eligible>;
  passive;
  poll-interval seconds;
  priority number;
  retransmit-interval seconds;
  te-metric metric;
  topology (ipv4-multicast | name) {
    metric metric;
  }
  transit-delay seconds;
  transmit-interval seconds;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols (ospf | ospf3) **area** *area-id*],
[edit logical-systems *logical-system-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) **area** *area-id*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf | ospf3) **area** *area-id*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) **area** *area-id*],
[edit protocols (ospf | ospf3) **area** *area-id*],
[edit protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) **area** *area-id*],

[edit routing-instances *routing-instance-name* protocols (ospf | ospf3) **area** *area-id*],
[edit routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) **area** *area-id*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable OSPF routing on a routing device interface.

You must include at least one **interface** statement in the configuration to enable OSPF on the routing device.

Options ***interface-name***—Name of the interface. Specify the interface by IP address or interface name for OSPFv2, or only the interface name for OSPFv3. Using both the interface name and IP address of the same interface produces an invalid configuration. To configure all interfaces, you can specify **all**. Specifying a particular interface and **all** produces an invalid configuration. For details about specifying interfaces, see interface naming in the *Junos OS Network Interfaces Configuration Guide*.

NOTE: For nonbroadcast interfaces, specify the IP address of the nonbroadcast interface as *interface-name*.

The remaining statements are explained separately.

NOTE: You cannot run both OSPF and ethernet-tcc encapsulation between two routing devices.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- neighbor
- Minimum OSPF Configuration
- Configuring Multitopology Routing in OSPF
- Configuring Multiple Address Families for OSPFv3

interface (Routing Options)

Syntax

```
interface interface-names {
    maximum-bandwidth bps;
    no-qos-adjust;
    reverse-oif-mapping {
        no-qos-adjust;
    }
    subscriber-leave-timer seconds;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Define the maximum bandwidth for an interface on which you want to apply bandwidth management.

Options

interface-name—Names of the physical or logical interface. For details about specifying interfaces, see the *Junos OS Network Interfaces Configuration Guide*.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

interface (Multicast via Static Routes)

Syntax

```
interface interface-names {
  maximum-bandwidth bps;
  no-qos-adjust;
  reverse-oif-mapping {
    no-qos-adjust;
  }
  subscriber-leave-timer seconds;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options **multicast**],
[edit logical-systems *logical-system-name* routing-options **multicast**],
[edit routing-instances *routing-instance-name* routing-options **multicast**],
[edit routing-options **multicast**]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Enable multicast traffic on an interface.

NOTE: You cannot enable multicast traffic on an interface using the enable statement and configure PIM on the same interface simultaneously.

Options

interface-name—Name of the interface on which to enable multicast traffic. Specify the ***interface-name*** to enable multicast traffic on the interface.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Enabling Multicast Forwarding Without PIM

interface-routes

Syntax

```
interface-routes {
  family (inet | inet6) {
    export {
      lan;
      point-to-point;
    }
  }
  rib-group group-name;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Associate a routing table group with the routing device's interfaces and specify routing table groups into which interface routes are imported.

Options

inet—Specify the IPv4 address family.

inet6—Specify the IPv6 address family.

lan—Export LAN routes.

point-to-point—Export point-to-point routes.

The remaining statement is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- passive
- Configuring How Interface Routes Are Imported into Routing Tables

interface-type

Syntax interface-type (nbma | p2mp | p2p);

Hierarchy Level
[edit logical-systems *logical-system-name* protocols (ospf | ospf3) area *area-id* **interface** *interface-name*],
[edit logical-systems *logical-system-name* protocols ospf3 realm (ipv4-multicast | ipv4-unicast | ipv6-multicast) area *area-id* interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 realm (ipv4-multicast | ipv4-unicast | ipv6-multicast) area *area-id* interface *interface-name*],
[edit protocols (ospf | ospf3) area *area-id* **interface** *interface-name*],
[edit protocols ospf3 realm (ipv4-multicast | ipv4-unicast | ipv6-multicast) area *area-id* interface *interface-name*]
[edit routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* **interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols ospf3 realm (ipv4-multicast | ipv4-unicast | ipv6-multicast) area *area-id* interface *interface-name*],

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the type of interface.

By default, the software chooses the correct interface type based on the type of physical interface. Therefore, you should never have to set the interface type. The exception to this is for NBMA interfaces, which default to an interface type of point-to-multipoint. To have these interfaces explicitly run in NBMA mode, configure the **nbma** interface type, using the IP address of the local ATM interface.

A point-to-point interface can be an Ethernet interface without a subnet. For more information about configuring interfaces, see the *Junos OS Network Interfaces Configuration Guide*.

Default The software chooses the correct interface type based on the type of physical interface.

Options **nbma** (OSPFv2 only)—Nonbroadcast multiaccess (NBMA) interface.

p2mp (OSPFv2 only)—Point-to-multipoint interface.

p2p—Point-to-point interface.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring OSPF on Interfaces

ipv4-multicast

Syntax ipv4-multicast;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols isis **topologies**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis **topologies**],
[edit protocols isis **topologies**],
[edit routing-instances *routing-instance-name* protocols isis **topologies**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure alternate IPv4 multicast topologies.

Default Multicast topologies are disabled.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring IS-IS Multicast Topologies

ipv4-multicast-metric

Syntax ipv4-multicast-metric *metric*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols isis interface *interface-name*level *level-number*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis interface *interface-name* level *level-number*],
[edit protocols isis interface *interface-name* level *level-number*],
[edit routing-instances *routing-instance-name* protocols isis interface *interface-name* level *level-number*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the multicast topology metric value for the level.

Options *metric*—Metric value.
Range: 0 through 16,777,215

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring IS-IS Multicast Topologies

ipv6-multicast

Syntax ipv6-multicast;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols isis **topologies**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis **topologies**],
[edit protocols isis **topologies**],
[edit routing-instances *routing-instance-name* protocols isis **topologies**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure alternate IPv6 multicast topologies.

Default Multicast topologies are disabled.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring IS-IS Multicast Topologies

ipv6-multicast-metric

Syntax ipv6-multicast-metric *metric*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols isis interface *interface-name* level *level-number*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis interface *interface-name* level *level-number*],
[edit protocols isis interface *interface-name* level *level-number*],
[edit routing-instances *routing-instance-name* protocols isis interface *interface-name* level *level-number*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the IPv6 alternate multicast topology metric value for the level.

Options *metric*—Metric value.
Range: 0 through 16,777,215

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring IS-IS Multicast Topologies

ipv6-unicast

Syntax ipv6-unicast;

Hierarchy Level [edit logical-systems *logical-system-name* protocols isis **topologies**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis **topologies**],
[edit protocols isis **topologies**],
[edit routing-instances *routing-instance-name* protocols isis **topologies**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure alternate IPv6 unicast topologies.

Default IPv6 unicast topologies are disabled.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring IS-IS IPv6 Unicast Topologies

ipv6-unicast-metric

Syntax ipv6-unicast-metric *metric*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols isis interface *interface-name* level *level-number*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis interface *interface-name* level *level-number*],
[edit protocols isis interface *interface-name* level *level-number*],
[edit routing-instances *routing-instance-name* protocols isis interface *interface-name* level *level-number*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the IPv6 unicast topology metric value for the level.

Options *metric*—Metric value.
Range: 0 through 16,777,215

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring IS-IS IPv6 Unicast Topologies

isis

Syntax isis { ... }

Hierarchy Level [edit logical-systems *logical-system-name* protocols],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols],
[edit protocols],
[edit routing-instances *routing-instance-name* protocols]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable IS-IS routing on the routing device or for a routing instance.

The **isis** statement is the one statement you must include in the configuration to run IS-IS on the routing device or in a routing instance.

Default IS-IS is disabled on the routing device.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Minimum IS-IS Configuration

keep

Syntax keep (all | none);

Hierarchy Level
[edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name* **neighbor** *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],
[edit protocols bgp],
[edit protocols bgp **group** *group-name*],
[edit protocols bgp group *group-name* **neighbor** *address*],
[edit routing-instances *routing-instance-name* protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify whether routes learned from a BGP peer are retained in the routing table even if they contain an AS number that was exported from the local AS.

Default If you do not include this statement, most routes are retained in the routing table.

Options **all**—Retain all routes.

none—Retain none of the routes. When **keep none** is configured for the BGP session and the inbound policy changes, the Junos OS forces readvertisement of the full set of routes advertised by the peer.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Applying Policies to BGP Routes

labeled-unicast

Syntax

```
labeled-unicast {
  accepted-prefix-limit {
    maximum number;
    teardown <percentage> <idle-timeout (forever | minutes)>;
  }
  aggregate-label {
    community community-name;
  }
  explicit-null {
    connected-only;
  }
  prefix-limit {
    maximum number;
    teardown <percentage> <idle-timeout (forever | minutes)>;
  }
  resolve-vpn;
  rib inet.3;
  rib-group group-name;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols bgp **family** (inet | inet6)],
[edit logical-systems *logical-system-name* protocols bgp group *group-name* **family** (inet | inet6)],
[edit logical-systems *logical-system-name* protocols bgp group *group-name* neighbor *address* **family** (inet | inet6)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **family** (inet | inet6)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **family** (inet | inet6)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address* **family** (inet | inet6)],
[edit protocols bgp **family** (inet | inet6)],
[edit protocols bgp group *group-name* **family** (inet | inet6)],
[edit protocols bgp group *group-name* neighbor *address* **family** (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols bgp **family** (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **family** (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address* **family** (inet | inet6)]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure the family type to be labeled-unicast.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Enabling Multiprotocol BGP

level (Global IS-IS)

Syntax

```
level level-number {
    authentication-key key;
    authentication-type type;
    external-preference preference;
    no-csnp-authentication;
    no-hello-authentication;
    no-psnp-authentication;
    preference preference;
    wide-metrics-only;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols **isis**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **isis**],
[edit protocols **isis**],
[edit routing-instances *routing-instance-name* protocols **isis**]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure the global-level properties.

Options

level-number—IS-IS level number.

Values: 1 or 2

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Preference Values for IS-IS Routes

link-protection

Syntax link-protection;

Hierarchy Level [edit logical-systems *logical-system-name* protocols isis interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis interface *interface-name*],
[edit protocols isis interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols isis interface *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable link protection on the specified IS-IS interface. The Junos OS creates a backup loop-free alternate path to the primary next hop for all destination routes that traverse the protected interface.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **node-link-protection on page 1641**
- Configuring Loop-Free Alternate Routes for IS-IS

local-address

Syntax local-address *address*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name* **neighbor** *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],
[edit protocols bgp],
[edit protocols bgp **group** *group-name*],
[edit protocols bgp group *group-name* **neighbor** *address*],
[edit routing-instances *routing-instance-name* protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the address of the local end of a BGP session. This address is used to accept incoming connections to the peer and to establish connections to the remote peer. When none of the operational interfaces are configured with the specified local address, a session with a BGP peer is placed in the idle state.

Default If you do not configure a local address, BGP uses the routing device's source address selection rules to set the local address. For more information, see the *Junos OS Network Interfaces Configuration Guide*.

Options ***address***—IPv6 or IPv4 address of the local end of the connection.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **router-id on page 1697**
- Enabling BGP

local-address

Syntax local-address *address*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast **backup-pe-group** *group-name*],
[edit logical-systems *logical-system-name* routing-options multicast **backup-pe-group** *group-name*],
[edit routing-instances *routing-instance-name* routing-options multicast **backup-pe-group** *group-name*],
[edit routing-options multicast **backup-pe-group** *group-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the address of the local PE for ingress PE redundancy when point-to-multipoint LSPs are used for multicast distribution.

Options ***address***—Address of local PEs in the backup group.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Ingress PE Redundancy

local-as

Syntax local-as *autonomous-system* <loops *number*> <private | alias> <no-prepend-global-as>;

Hierarchy Level [edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name* neighbor *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],
[edit protocols bgp],
[edit protocols bgp **group** *group-name*],
[edit protocols bgp group *group-name* **neighbor** *address*],
[edit routing-instances *routing-instance-name* protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set the local AS number.

The autonomous system (AS) numeric range in plain-number format provides BGP support for 4-byte AS numbers, as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*.

You can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: *<16-bit high-order value in decimal>.<16-bit low-order value in decimal>*. For example, the 4-byte AS number of 65546 in plain-number format is represented as 1.10 in the AS-dot notation format.

Options **alias**—(Optional) Configure the local AS as an alias of the global AS number configured for the router at the **[edit routing-options]** hierarchy level. As a result, a BGP peer considers any local AS to which it is assigned as equivalent to the primary AS number configured for the routing device. When you use the **alias** option, only the AS (global or local) used to establish the BGP session is prepended in the AS path sent to the BGP neighbor.

autonomous-system—AS number.

Range: 1 through 4,294,967,295 ($2^{32} - 1$) in plain-number format

Range: 0.0 through 65535.65535 in AS-dot notation format

loops *number*—(Optional) Specify the maximum number of times that the local AS number can appear in an AS path received from a BGP peer. For ***number***, include a value from 1 through 10.

no-prepend-global-as—(Optional) Specify to strip the global AS and to prepend only the local AS in AS paths sent to external peers.

private—(Optional) Configure to use the local AS only during the establishment of the BGP session with a BGP neighbor but to hide it in the AS path sent to external BGP peers. Only the global AS is included in the AS path sent to external peers.

NOTE: The private and alias options are mutually exclusive. You cannot configure both options with the same local-as statement.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **autonomous-system on page 1487**
- **family on page 1536**
- Configuring a Local AS for EBGP Sessions

local-interface

Syntax
local-interface *interface-name*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name* **neighbor** *ipv6-link-local-address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *ipv6-link-local-address*],
[edit protocols bgp group *group-name* **neighbor** *ipv6-link-local-address*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *ipv6-link-local-address*]

Release Information
Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Specify the interface name of the peer for IPv6 peering using link-local addresses. This peer is link-local in scope.

Options
interface-name—Interface name of the EBGP IPv6 peer.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring EBGP Peering Using IPv6 Link-Local Addresses

local-preference

Syntax local-preference *local-preference*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name* **neighbor** *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],
[edit protocols bgp],
[edit protocols bgp **group** *group-name*],
[edit protocols bgp group *group-name* **neighbor** *address*],
[edit routing-instances *routing-instance-name* protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Modify the value of the **LOCAL_PREF** path attribute, which is a metric used by IBGP sessions to indicate the degree of preference for an external route. The route with the highest local preference value is preferred.

The **LOCAL_PREF** path attribute always is advertised to internal BGP peers and to neighboring confederations. It is never advertised to external BGP peers.

Default If you omit this statement, the **LOCAL_PREF** path attribute, if present, is not modified.

Options *local-preference*—Preference to assign to routes learned from BGP or from the group or peer.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

Default: If the **LOCAL_PREF** path attribute is present, do not modify its value. If a BGP route is received without a **LOCAL_PREF** attribute, the route is handled locally (it is stored in the routing table and advertised by BGP) as if it were received with a **LOCAL_PREF** value of 100. By default, non-BGP routes that are advertised by BGP are advertised with a **LOCAL_PREF** value of 100.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **preference on page 1660**
- Configuring the Local Preference Value for BGP Routes

log-updown

Syntax log-updown;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name* **neighbor** *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],
[edit protocols bgp],
[edit protocols bgp **group** *group-name*],
[edit protocols bgp group *group-name* **neighbor** *address*],
[edit routing-instances *routing-instance-name* protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Log a message whenever a BGP peer makes a state transition. Messages are logged using the system logging mechanism located at the **[edit system syslog]** hierarchy level.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **traceoptions on page 1718**
- Configuring System Logging of BGP Peer State Transitions
- *Junos OS System Basics Configuration Guide*

loose-authentication-check

Syntax loose-authentication-check;

Hierarchy Level [edit logical-systems *logical-system-name* protocols **isis**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **isis**],
[edit protocols **isis**],
[edit routing-instances *routing-instance-name* protocols **isis**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Allow the use of MD5 authentication without requiring network-wide deployment.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Enabling Authentication for IS-IS Without Network-Wide Deployment

lsp-interval

Syntax lsp-interval *milliseconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols isis **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis **interface** *interface-name*],
[edit protocols isis **interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols isis **interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the link-state PDU interval time.

Options ***milliseconds***—Number of milliseconds between the sending of link-state PDUs. Specifying a value of 0 blocks all link-state PDU transmission.

Range: 0 through 1000 milliseconds

Default: 100 milliseconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring the Transmission Frequency for Link-State PDUs on IS-IS Interfaces

lsp-lifetime

Syntax lsp-lifetime *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols **isis**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **isis**],
[edit protocols **isis**],
[edit routing-instances *routing-instance-name* protocols **isis**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify how long a link-state PDU originating from the routing device should persist in the network. The routing device sends link-state PDUs often enough so that the link-state PDU lifetime never expires.

Options ***seconds***—link-state PDU lifetime, in seconds.

Range: 350 through 65,535 seconds

Default: 1200 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Link-State PDU Lifetime for IS-IS

lsp-metric-into-summary

Syntax lsp-metric-into-summary;

Hierarchy Level [edit logical-systems *logical-system-name* protocols (ospf | ospf3) **traffic-engineering shortcuts**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf | ospf3) **traffic-engineering shortcuts**],
[edit protocols (ospf | ospf3) **traffic-engineering shortcuts**],
[edit routing-instances *routing-instance-name* protocols (ospf | ospf3) **traffic-engineering shortcuts**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Advertise the LSP metric in summary LSAs.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Enabling OSPF Traffic Engineering Support

martians

Syntax

```
martians {
  destination-prefix match-type <allow>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options **rib** *routing-table-name*],
[edit logical-systems *logical-system-name* routing-options],
[edit logical-systems *logical-system-name* routing-options **rib** *routing-table-name*],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options **rib** *routing-table-name*],
[edit routing-options],
[edit routing-options **rib** *routing-table-name*]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure martian addresses.

Options

allow—(Optional) Explicitly allow a subset of a range of addresses that has been disallowed.

destination-prefix—Destination route you are configuring:

- ***destination-prefix/prefix-length***—***destination-prefix*** is the network portion of the IP address, and ***prefix-length*** is the destination prefix length.
- **default**—Default route to use when routing packets do not match a network or host in the routing table. This is equivalent to specifying the IP address **0.0.0.0/0**.

match-type—Criteria that the destination must match:

- **exact**—Exactly match the route's mask length.
- **longer**—The route's mask length is greater than the specified mask length.
- **orlonger**—The route's mask length is equal to or greater than the specified mask length.
- **through *destination-prefix***—The route matches the first prefix, the route matches the second prefix for the number of bits in the route, and the number of bits in the route is less than or equal to the number of bits in the second prefix.
- **upto *prefix-length***—The route's mask length falls between the two destination prefix lengths, inclusive.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Martian Addresses

max-areas

Syntax max-areas *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols **isis**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **isis**]
[edit protocols **isis**],
[edit routing-instances *routing-instance-name* protocols **isis**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Modify the maximum number of IS-IS areas advertised.

Options ***number***—Maximum number of areas to include in the IS-IS hello (IIH) PDUs and link-state PDUs.

Range: 3 through 36

Default: 3

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Limiting the Number of Advertised IS-IS Areas

maximum-bandwidth

Syntax maximum-bandwidth *bps*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-options multicast **interface** *interface-name*],
[edit routing-instances *routing-instance-name* routing-options multicast **interface** *interface-name*],
[edit routing-options multicast **interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the multicast bandwidth for the interface.

Options ***bps***—Bandwidth rate, in bits per second, for the multicast interface.

Range: 0 through any amount of bandwidth

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Defining Interface Bandwidth Maximums

maximum-paths

Syntax maximum-paths *path-limit* <log-interval *seconds*> <log-only | threshold *value*>;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a limit for the number of routes installed in a routing table based upon the route path.

Options

log-interval *seconds*—(Optional) Minimum time interval (in seconds) between log messages.

Range: 5 through 86,400

log-only—(Optional) Sets the route limit as an advisory limit. An advisory limit triggers only a warning, and additional routes are not rejected.

path-limit—Maximum number of routes. If this limit is reached, a warning is triggered and additional routes are rejected.

Range: 1 through 4,294,967,295 ($2^{32} - 1$)

Default: No default

threshold *value*—(Optional) Percentage of the maximum number of routes that starts triggering warning. You can configure a percentage of the ***path-limit*** value that starts triggering the warnings.

Range: 1 through 100

NOTE: When the number or routes reaches the **threshold** value, routes are still installed into the routing table while warning messages are sent. When the number or routes reaches the ***path-limit*** value, then additional routes are rejected.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Route Limits for Routing Tables

maximum-prefixes

Syntax maximum-prefixes *prefix-limit* <log-interval *seconds*> <log-only | threshold *value*>;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a limit for the number of routes installed in a routing table based upon the route prefix.

Options

log-interval *seconds*—(Optional) Minimum time interval (in seconds) between log messages.

Range: 5 through 86,400

log-only—(Optional) Sets the prefix limit as an advisory limit. An advisory limit triggers only a warning, and additional routes are not rejected.

prefix-limit—Maximum number of route prefixes. If this limit is reached, a warning is triggered and any additional routes are rejected.

Range: 1 through 4,294,967,295

Default: No default

threshold *value*—(Optional) Percentage of the maximum number of prefixes that starts triggering warning. You can configure a percentage of the ***prefix-limit*** value that starts triggering the warnings.

Range: 1 through 100

NOTE: When the number or routes reaches the threshold value, routes are still installed into the routing table while warning messages are sent. When the number or routes reaches the *prefix-limit* value, then additional routes are rejected.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Route Limits for Routing Tables

med-igp-update-interval

Syntax med-igp-update-interval *minutes*;

Hierarchy Level [edit routing-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a timer for how long to delay updates for the multiple-exit discriminator (MED) path attribute for BGP groups and peers configured with the **metric-out igp *offset* delay-med-update** statement. The timer delays MED updates for the interval configured unless the MED is lower than the previously advertised attribute or another attribute associated with the route has changed or if the BGP peer is responding to a refresh route request.

Options ***minutes***—Interval to delay MED updates.

Default: 10 minutes

Range: 10 through 600

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **metric-out on page 1615**
- Delaying Updates of the MED Path Attribute for BGP

mesh-group

Syntax mesh-group (blocked | *value*);

Hierarchy Level [edit logical-systems *logical-system-name* protocols isis **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis **interface** *interface-name*],
[edit protocols isis **interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols isis **interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure an interface to be part of a mesh group, which is a set of fully connected nodes.

Options **blocked**—Configure the interface so that it does not flood link-state PDU packets.

value—Number that identifies the mesh group.

Range: 1 through 4,294,967,295 (2^{32} – 1; 32 bits are allocated to identify a mesh group)

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Mesh Groups of IS-IS Interfaces

message-size

Syntax message-size *number*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols **rip**],
[edit logical-systems *logical-system-name* protocols rip group *group-name* **neighbor** *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **rip**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip group *group-name* **neighbor** *neighbor-name*],
[edit protocols **rip**],
[edit protocols rip group *group-name* **neighbor** *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols **rip**],
[edit routing-instances *routing-instance-name* protocols rip group *group-name* **neighbor** *neighbor-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the number of route entries to be included in every RIP update message. To ensure interoperability with other vendors' equipment, use the standard of 25 route entries per message.

Options ***number***—Number of route entries per update message.

Range: 25 through 255 entries

Default: 25 entries

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring the Number of Route Entries in RIP Update Messages

metric

Syntax metric *metric*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols isis interface *interface-name* level *level-number*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis interface *interface-name* level *level-number*],
[edit protocols isis interface *interface-name* level *level-number*],
[edit routing-instances *routing-instance-name* protocols isis interface *interface-name* level *level-number*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the metric value for the level.

Options *metric*—Metric value.

Range: 1 through 63, or 1 through 16,777,215 (if you have configured wide metrics)

Default: 10 (for all interfaces except **lo0**), 0 (for the **lo0** interface)

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- te-metric
- **wide-metrics-only on page 1745**
- Configuring Levels on IS-IS Interfaces

metric

Syntax metric *metric*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols (ospf | ospf3) area *area-id* interface *interface-name*],
[edit logical-systems *logical-system-name* protocols ospf area *area-id* interface *interface-name* topology (ipv4-multicast | *name*)],
[edit logical-systems *logical-system-name* protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf area *area-id* sham-link-remote],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf area *area-id* interface *interface-name* topology (ipv4-multicast | *name*)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* interface *interface-name*],
[edit protocols (ospf | ospf3) area *area-id* interface *interface-name*],
[edit protocols ospf area *area-id* interface *interface-name* topology (ipv4-multicast | *name*)],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols ospf area *area-id* sham-link-remote],
[edit routing-instances *routing-instance-name* protocols ospf area *area-id* interface *interface-name* topology (ipv4-multicast | *name*)],
[edit routing-instances *routing-instance-name* protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* interface *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the cost of an OSPF interface. The cost is a routing metric that is used in the link-state calculation.

To set the cost of routes exported into OSPF, configure the appropriate routing policy.

Options *metric*—Cost of the route.

Range: 1 through 65,535

Default: By default, the cost of an OSPF route is calculated by dividing the reference-bandwidth value by the bandwidth of the physical interface. Any specific value you configure for the **metric** overrides the default behavior of using the reference-bandwidth value to calculate the cost of route for that interface.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **bandwidth-based-metrics on page 1491**
- **reference-bandwidth on page 1677**

- Configuring the Metric Value for OSPF Interfaces
- Configuring OSPF Sham Links
- Configuring Multitopology Routing in OSPF

metric (Aggregate, Generated, or Static Route)

Syntax (metric | metric2 | metric3 | metric4) *metric* <type *type*>;

Hierarchy Level [edit logical-systems *logical-system-name* routing-options (**aggregate** | **generate** | **static**) (defaults | route)],
[edit routing-options (**aggregate** | **generate** | **static**) (defaults | route)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Metric value for an aggregate, generated, or static route. You can specify up to four metric values, starting with **metric** (for the first metric value) and continuing with **metric2**, **metric3**, and **metric4**.

Options *metric*—Metric value.
Range: 0 through 4,294,967,295 ($2^{32} - 1$)

type *type*—(Optional) Type of route.
Range: 1 through 16

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **aggregate on page 1470**
- **generate on page 1543**
- **static on page 1708**
- Configuring Static Route Options
- Configuring Aggregate Route Options
- Configuring Generated Route Options

metric-in

Syntax metric-in *metric*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols **rip**],
[edit logical-systems *logical-system-name* protocols rip group *group-name* **neighbor** *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **rip**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip group *group-name* **neighbor** *neighbor-name*],
[edit protocols **rip**],
[edit protocols rip group *group-name* **neighbor** *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols **rip**],
[edit routing-instances *routing-instance-name* protocols rip group *group-name* **neighbor** *neighbor-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the metric to add to incoming routes when advertising into RIP routes that were learned from other protocols. Use this statement to configure the routing device to prefer RIP routes learned through a specific neighbor.

Options *metric*—Metric value.

Range: 1 through 16

Default: 1

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring the Metric Value Added to Imported RIP Routes

metric-in

Syntax metric-in *metric*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols ripng],
[edit logical-systems *logical-system-name* protocols ripng group *group-name* neighbor *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng group *group-name* neighbor *neighbor-name*],
[edit protocols ripng],
[edit protocols ripng group *group-name* neighbor *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols ripng],
[edit routing-instances *routing-instance-name* protocols ripng group *group-name* neighbor *neighbor-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the metric to add to incoming routes when advertising into RIPng routes that were learned from other protocols. Use this statement to configure the routing device to prefer RIPng routes learned through a specific neighbor.

Options *metric*—Metric value.

Range: 1 through 16

Default: 1

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring the Metric Value Added to Imported RIPng Routes

metric-out

Syntax metric-out (*metric* | minimum-igp *offset* | igp (delay-med-update | *offset*);

Hierarchy Level [edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name* **neighbor** *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],
[edit protocols bgp],
[edit protocols bgp **group** *group-name*],
[edit protocols bgp group *group-name* **neighbor** *address*],
[edit routing-instances *routing-instance-name* protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Metric for all routes sent using the multiple exit discriminator (MED, or **MULTI_EXIT_DISC**) path attribute in update messages. This path attribute is used to discriminate among multiple exit points to a neighboring AS. If all other factors are equal, the exit point with the lowest metric is preferred.

You can specify a constant metric value by including the ***metric*** option. For configurations in which a BGP peer sends third-party next hops that require the local system to perform next-hop resolution—IBGP configurations, configurations within confederation peers, or EBGP configurations that include the **multihop** command—you can specify a variable metric by including the **minimum-igp** or **igp** option.

You can increase or decrease the variable metric calculated from the IGP metric (either from the **igp** or **igp-minimum** statement) by specifying a value for ***offset***. The metric is increased by specifying a positive value for ***offset***, and decreased by specifying a negative value for ***offset***.

You can specify for a BGP group or peer not to advertise updates for the MED path attributes used to calculate IGP costs for BGP next hops unless the MED is lower. You can also configure an interval to delay when MED updates are sent by including the **med-igp-update-interval** ***minutes*** at the **[edit routing-options]** hierarchy level.

Options

delay-med-update—Specify for a BGP group or peer configured with the **metric-out igp** statement not to advertise MED updates when the value worsens, that is, unless the value is lower.

NOTE: You cannot configure delay-med–update statement at the global BGP level.

igp—Set the metric to the most recent metric value calculated in the IGP to get to the BGP next hop.

metric—Primary metric on all routes sent to peers.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

Default: No metric is sent.

minimum-igp—Set the metric to the minimum metric value calculated in the IGP to get to the BGP next hop. If a newly calculated metric is greater than the minimum metric value, the metric value remains unchanged. If a newly calculated metric is lower, the metric value is lowered to that value.

offset—(Optional) Increases or decreases the metric by this value.

Range: -2^{31} through $2^{31} - 1$

Default: None

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **med-igp-update-interval on page 1607**
- Configuring the MED in BGP Updates

metric-out

Syntax metric-out *metric*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit protocols rip group *group-name* neighbor *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the metric value to add to routes transmitted to the neighbor. Use this statement to control how other routing devices prefer RIP routes sent from this neighbor.

Options ***metric***—Metric value.

Range: 1 through 16

Default: 1

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Group-Specific RIP Properties

metric-out

Syntax metric-out *metric*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols ripng group *group-name* neighbor *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng group *group-name* neighbor *neighbor-name*],
[edit protocols ripng group *group-name* neighbor *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols ripng group *group-name* neighbor *neighbor-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the metric value to add to routes transmitted to the neighbor. Use this statement to control how other routing devices prefer RIPng routes sent from this neighbor.

Options ***metric***—Metric value.

Range: 1 through 16

Default: 1

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Group-Specific RIPng Properties

metric-type

Syntax metric-type *type*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols (ospf | ospf3) area *area-id* **nssa default-lsa**],
[edit logical-systems *logical-system-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)] area *area-id* nssa**default-lsa**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* **nssa default-lsa**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)] area *area-id* nssa **default-lsa**],
[edit protocols (ospf | ospf3) area *area-id* **nssa default-lsa**],
[edit protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)] area *area-id* nssa **default-lsa**],
[edit routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* **nssa default-lsa**],
[edit routing-instances *routing-instances* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)] area *area-id* nssa **default-lsa**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the external metric type for the default LSA.

Options *type*—Metric type: **1** or **2**

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring OSPF Areas

mtu-discovery

Syntax mtu-discovery;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name* **neighbor** *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],
[edit protocols bgp],
[edit protocols bgp **group** *group-name*],
[edit protocols bgp group *group-name* **neighbor** *address*],
[edit routing-instances *routing-instance-name* protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure TCP path maximum transmission unit (MTU) discovery. MTU discovery improves convergence times for IBGP sessions.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring MTU Discovery for BGP Sessions

multicast

Syntax

```
multicast {
  forwarding-cache {
    threshold suppress value <reuse value>;
  }
  interface interface-name {
    enable;
  }
  scope scope-name {
    interface [ interface-names ];
    prefix destination-prefix;
  }
  ssm-groups {
    address;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure generic multicast properties.

NOTE: You cannot apply a scoping policy to a specific routing instance. All scoping policies are applied to all routing instances. However, you can apply the scope statement to a specific routing instance.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- (**indirect-next-hop on page 1573** | no-indirect-next-hop)
- Configuring Multicast Scoping
- Configuring Additional Source-Specific Multicast Groups
- *Junos OS Multicast Configuration Guide*

multihop

Syntax

```
multihop {
    no-nexthop-change;
    ttl-value;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name* **neighbor** *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],
[edit protocols bgp],
[edit protocols bgp **group** *group-name*],
[edit protocols bgp group *group-name* **neighbor** *address*],
[edit routing-instances *routing-instance-name* protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure an EBGP multihop session.

External confederation peering is a special case that allows unconnected third-party next hops. You do not need to configure multihop sessions explicitly in this particular case; multihop behavior is implied.

If you have confederation external BGP peer-to-loopback addresses, you still need the multihop configuration.

Default

If you omit this statement, all EBGP peers are assumed to be directly connected (that is, you are establishing a nonmultihop, or "regular," BGP session), and the default time-to-live (TTL) value is 1.

Options

no-nexthop-change—Specify not to change the BGP next-hop value; for route advertisements, specify the **no-nexthop-self** option.

ttl-value—Configure the maximum TTL value for the TTL in the IP header of BGP packets.

Range: 1 through 255

Default: 64 (for multihop EBGP sessions, confederations, and IBGP sessions)

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring EBGP Multihop Sessions

multipath

Syntax

```
multipath {
  multiple-as;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name* **neighbor** *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],
[edit protocols bgp **group** *group-name*],
[edit protocols bgp group *group-name* **neighbor** *address*],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Allow load sharing among multiple EBGP paths and multiple IBGP paths.

Options

multiple-as—Disable the default check requiring that paths accepted by BGP multipath must have the same neighboring AS.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Selecting Multiple Equal-Cost Active Paths

neighbor

Syntax

```
neighbor address {
    accept-remote-nexthop;
    advertise-external <conditional>;
    advertise-inactive;
    (advertise-peer-as | no-advertise-peer-as);
    as-override;
    authentication-algorithm algorithm;
    authentication-key key;
    authentication-key-chain key-chain;
    cluster cluster-identifier;
    damping;
    description text-description;
    export [ policy-names ];
    family {
        (inet | inet6 | inet-mvpn | inet6-mpvn | inet-vpn | inet6-vpn | iso-vpn | l2-vpn) {
            (any | flow | multicast | unicast | signaling) {
                accepted-prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                }
                prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                }
                rib-group group-name;
            }
            flow {
                no-validate policy-name;
            }
            labeled-unicast {
                accepted-prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                }
                aggregate-label {
                    community community-name:
                }
                explicit-null {
                    connected-only;
                }
                prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | minutes)>;
                }
                resolve-vpn;
                rib inet.3;
                rib-group group-name;
            }
        }
        route-target {
            advertise-default;
            external-paths number;
```

```
        accepted-prefix-limit {
          maximum number;
          teardown <percentage> <idle-timeout (forever | minutes)>;
        }
        prefix-limit {
          maximum number;
          teardown <percentage> <idle-timeout (forever | minutes)>;
        }
      }
      signaling {
        prefix-limit {
          maximum number;
          teardown <percentage> <idle-timeout (forever | minutes)>;
        }
      }
    }
    graceful-restart {
      disable;
      restart-time seconds;
      stale-routes-time seconds;
    }
    hold-time seconds;
    import [ policy-names ];
    ipsec-sa ipsec-sa;
    keep (all | none);
    local-address address;
    local-as autonomous-system <private>;
    local-interface interface-name;
    local-preference preference;
    log-updown;
    metric-out (metric | minimum-igp <offset> | igp <offset>);
    mtu-discovery;
    multihop <ttl-value>;
    multipath {
      multiple-as;
    }
    no-aggregator-id;
    no-client-reflect;
    out-delay seconds;
    passive;
    peer-as autonomous-system;
    preference preference;
    tcp-mss segment-size;
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>;
      flag flag <flag-modifier> <disable>;
    }
    vpn-apply-export;
  }
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Explicitly configure a neighbor (peer). To configure multiple BGP peers, include multiple **neighbor** statements.

By default, the peer's options are identical to those of the group. You can override these options by including peer-specific option statements within the **neighbor** statement.

The **neighbor** statement is one of the statements you can include in the configuration to define a minimal BGP configuration on the routing device. (You can include an **allow all** statement in place of a **neighbor** statement.)

Options ***address***—IPv6 or IPv4 address of a single peer.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Minimum BGP Configuration
- Configuring BGP Groups and Peers

neighbor

Syntax

```
neighbor neighbor-name {
  authentication-key password;
  authentication-type type;
  bfd-liveness-detection {
    authentication {
      algorithm algorithm-name;
      key-chain key-chain-name;
      loose-check;
    }
    detection-time {
      threshold milliseconds;
    }
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    transmit-interval {
      threshold milliseconds;
      minimum-interval milliseconds;
    }
    multiplier number;
    version (0 | 1 | automatic);
  }
  (check-zero | no-check-zero);
  import policy-name;
  message-size number;
  metric-in metric;
  metric-out metric;
  receive receive-options;
  route-timeout seconds;
  send send-options;
  update-interval seconds;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols rip **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip **group** *group-name*],
[edit protocols rip **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols rip **group** *group-name*]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure neighbor-specific RIP parameters, thereby overriding the defaults set for the routing device.

Options

neighbor-name—Name of an interface over which a routing device communicates to its neighbors.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Overview of RIP Neighbor Properties

neighbor

Syntax

```
neighbor neighbor-name {
  import [ policy-names ];
  metric-in metric;
  receive <none>;
  route-timeout seconds;
  send <none>;
  update-interval seconds;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols ripng **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng group *group-name*],
[edit protocols ripng **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols ripng group *group-name*]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure neighbor-specific RIPng parameters, thereby overriding the defaults set for the routing device.

Options

neighbor-name—Name of an interface over which a routing device communicates to its neighbors.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Overview of RIPng Neighbor Properties

no-adjacency-holddown

Syntax no-adjacency-holddown;

Hierarchy Level [edit logical-systems *logical-system-name* protocols isis],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis],
[edit protocols isis],
[edit routing-instances *routing-instance-name* protocols isis]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable the hold-down timer for IS-IS adjacencies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Quicker Advertisement of IS-IS Adjacency State Changes

no-aggregator-id

Syntax no-aggregator-id;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name* neighbor *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],
[edit protocols bgp],
[edit protocols bgp **group** *group-name*],
[edit protocols bgp group *group-name* **neighbor** *address*],
[edit routing-instances *routing-instance-name* protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set the router ID in the BGP aggregator path attribute to zero. (This is one of the path attributes included in BGP update messages.) Doing this prevents different routing devices within an AS from creating aggregate routes that contain different AS paths.

Default If you omit this statement, the router ID is included in the BGP aggregator path attribute.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Overview of BGP Messages
- Controlling BGP Route Aggregation

no-authentication-check

Syntax no-authentication-check;

Hierarchy Level [edit logical-systems *logical-system-name* protocols isis],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis],
[edit protocols isis],
[edit routing-instances *routing-instance-name* protocols isis]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Generate authenticated packets and check the authentication on received packets, but do not reject packets that cannot be authenticated.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **csnp-interval on page 1516**
- **hello-authentication-type on page 1557**
- Configuring IS-IS Authentication

no-client-reflect

Syntax no-client-reflect;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name* **neighbor** *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],
[edit protocols bgp],
[edit protocols bgp **group** *group-name*],
[edit protocols bgp group *group-name* **neighbor** *address*],
[edit routing-instances *routing-instance-name* protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable intracluster route redistribution by the system acting as the route reflector. Include this statement when the client cluster is fully meshed to prevent the sending of redundant route advertisements.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **cluster on page 1513**
- Configuring BGP Route Reflection

no-csnp-authentication

Syntax no-csnp-authentication;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols isis **level** *level-number*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis **level** *level-number*],
[edit protocols isis **level** *level-number*],
[edit routing-instances *routing-instance-name* protocols isis **level** *level-number*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Suppress authentication check on complete sequence number PDU (CSNP) packets.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **csnp-interval on page 1516**
- Configuring IS-IS Authentication

no-eligible-backup

Syntax no-eligible-backup;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols isis interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis interface *interface-name*],
[edit protocols isis interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols isis interface *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Exclude the specified interface as a backup interface for IS-IS interfaces on which link protection or node-link protection is enabled.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **link-protection on page 1594**
- **node-link-protection on page 1641**
- Configuring Loop-Free Alternate Routes for IS-IS

no-hello-authentication

Syntax no-hello-authentication;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols isis **level** *level-number*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis **level** *level-number*],
[edit protocols isis **level** *level-number*],
[edit routing-instances *routing-instance-name* protocols isis **level** *level-number*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Suppress authentication check on complete sequence number hello packets.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **hello-authentication-type on page 1557**
- Configuring IS-IS Authentication

no-ipv4-multicast

Syntax no-ipv4-multicast;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols isis **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis **interface** *interface-name*],
[edit protocols isis **interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols isis **interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Exclude an interface from the IPv4 multicast topologies.

Default Multicast topologies are disabled.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring IS-IS Multicast Topologies

no-ipv4-routing

Syntax no-ipv4-routing;

Hierarchy Level [edit logical-systems *logical-system-name* protocols **isis**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **isis**],
[edit protocols **isis**],
[edit routing-instances *routing-instance-name* protocols **isis**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable IP version 4 (IPv4) routing.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Disabling IPv4 Routing for IS-IS

no-ipv6-multicast

Syntax no-ipv6-multicast;

Hierarchy Level [edit logical-systems *logical-system-name* protocols isis **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis **interface** *interface-name*],
[edit protocols isis **interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols isis **interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Exclude an interface from the IPv6 multicast topologies.

Default Multicast topologies are disabled.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring IS-IS Multicast Topologies

no-ipv6-routing

Syntax no-ipv6-routing;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols **isis**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **isis**],
[edit protocols **isis**],
[edit routing-instances *routing-instance-name* protocols **isis**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable IP version 6 (IPv6) routing.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Disabling IPv6 Routing for IS-IS

no-ipv6-unicast

Syntax no-ipv6-unicast;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols isis **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis **interface** *interface-name*],
[edit protocols isis **interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols isis **interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Exclude an interface from the IPv6 unicast topologies.

Default IPv6 unicast topologies are disabled.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring IS-IS IPv6 Unicast Topologies

no-nssa-abr

Syntax no-nssa-abr;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols (**ospf** | **ospf3**)],
[edit logical-systems *logical-system-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit protocols (**ospf** | **ospf3**)],
[edit protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)],
[edit routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable exporting Type 7 link-state advertisements into not-stubby-areas (NSSAs) for an autonomous system boundary router (ASBR) or an area border router (ABR).

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Disabling Export of LSAs into NSSAs Attached to ASBR ABRs

no-psnp-authentication

Syntax no-psnp-authentication;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols isis **level** *level-number*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis **level** *level-number*],
[edit protocols isis **level** *level-number*],
[edit routing-instances *routing-instance-name* protocols isis **level** *level-number*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Suppress authentication check on partial sequence number PDU (PSNP) packets.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring IS-IS Authentication

no-qos-adjust

Syntax no-qos-adjust;

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast **interface** *interface-name* **reverse-oif-mapping**],
[edit logical-systems *logical-system-name* routing-options multicast **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-options multicast **interface** *interface-name* **reverse-oif-mapping**],
[edit routing-instances *routing-instance-name* routing-options multicast **interface** *interface-name*],
[edit routing-instances *routing-instance-name* routing-options multicast **interface** *interface-name* **reverse-oif-mapping**],
[edit routing-options multicast **interface** *interface-name*],
[edit routing-options multicast **interface** *interface-name* **reverse-oif-mapping**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable hierarchical bandwidth adjustment for all subscriber interfaces that are identified by their MLD or IGMP request from a specific multicast interface.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Managing Subscriber Overcommitment

no-rfc-1583

Syntax no-rfc-1583;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols (**ospf** | **ospf3**)],
[edit logical-systems *logical-system-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit protocols (**ospf** | **ospf3**)],
[edit protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)],
[edit routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable compatibility with RFC 1583, *OSPF Version 2*. If the same external destination is advertised by AS boundary routers that belong to different OSPF areas, disabling compatibility with RFC 1583 can prevent routing loops.

Default Compatibility with RFC 1583 is enabled by default.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control-level—To add this statement to the configuration.

Related Documentation
- Disabling OSPFv2 Compatibility with RFC 1583

no-unicast-topology

Syntax no-unicast-topology;

Hierarchy Level [edit logical-systems *logical-system-name* protocols isis **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis **interface** *interface-name*],
[edit protocols isis **interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols isis **interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Exclude an interface from the IPv4 unicast topologies.

Default IPv4 unicast topologies are disabled.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring IS-IS Multicast Topologies

no-validate

Syntax no-validate *policy-name*;

Hierarchy Level [edit protocols bgp group *group-name* **family** (inet | inet flow)],
[edit protocols bgp group *group-name* neighbor *address* **family** (inet | inet flow)],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **family** (inet | inet flow)],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address* **family** (inet | inet flow)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Omits the flow route validation procedure after packets are accepted by a policy.

Options ***policy-name***—Import policy to match NLRI messages.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Enabling BGP to Carry Flow-Specification Routes

node-link-protection

Syntax node-ink-protection;

Hierarchy Level [edit logical-systems *logical-system-name* protocols isis interface *interface-name*],
[edit logical-routers *logical-router-name* routing-instances *routing-instance-name* protocols isis interface *interface-name*],
[edit protocols isis interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols isis interface *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable node-link protection on the specified IS-IS interface. The Junos OS creates an alternate loop-free path to the primary next hop for all destination routes that traverse a protected interface. This alternate path avoids the primary next-hop routing device altogether and establishes a path through a different routing device.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **link-protection on page 1594**
- Configuring Loop-Free Alternate Routes for IS-IS

nssa

Syntax

```
nssa {
  area-range network/mask-length <restrict> <exact> <override-metric metric>;
  default-lsa {
    default-metric metric;
    metric-type type;
    type-7;
  }
  (no-summaries | summaries);
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols (ospf | ospf3) area *area-id*],
[edit logical-systems *logical-system-name* protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit protocols (ospf | ospf3) area *area-id*],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id*],
[edit routing-instances *routing-instance-name* protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure a not-so-stubby area (NSSA). An NSSA allows external routes to be flooded within the area. These routes are then leaked into other areas.

You cannot configure an area as being both a stub area and an NSSA.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **stub on page 1710**
- Configuring OSPF Areas

options

Syntax

```
options {
  syslog (level level | upto level level);
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure the types of system logging messages sent about the routing protocols process to the system message logging file. These messages are also displayed on the system console. You can log messages at a particular level, or up to and including a particular level.

Options

level *level*—Severity of the message. It can be one or more of the following levels, in order of decreasing urgency:

- **alert**—Conditions that should be corrected immediately, such as a corrupted system database.
- **critical**—Critical conditions, such as hard drive errors.
- **debug**—Software debugging messages.
- **emergency**—Panic or other conditions that cause the system to become unusable.
- **error**—Standard error conditions.
- **info**—Informational messages.
- **notice**—Conditions that are not error conditions, but might warrant special handling.
- **warning**—System warning messages.

upto level *level*—Log all messages up to a particular level.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **syslog** in the *Junos OS System Basics Configuration Guide*
- Configuring System Logging for the Routing Protocol Process

ospf

Syntax ospf { ... }

Hierarchy Level
[edit logical-systems *logical-system-name* protocols],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols],
[edit protocols],
[edit routing-instances *routing-instance-name* protocols]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable OSPF routing on the routing device.

You must include the **ospf** statement to enable OSPF on the routing device.

Default OSPF is disabled on the routing device.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Minimum OSPF Configuration

ospf3

Syntax ospf3 { ... }

Hierarchy Level
[edit logical-systems *logical-system-name* protocols],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols],
[edit protocols],
[edit routing-instances *routing-instance-name* protocols]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable OSPFv3 routing on the routing device.

You must include the **ospf3** statement to enable OSPFv3.

Default OSPFv3 is disabled on the routing device.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Minimum OSPF Configuration

out-delay

Syntax out-delay *seconds*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name* **neighbor** *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],
[edit protocols bgp],
[edit protocols bgp **group** *group-name*],
[edit protocols bgp group *group-name* **neighbor** *address*],
[edit routing-instances *routing-instance-name* protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify how long a route must be present in the Junos OS routing table before it is exported to BGP. Use this time delay to help bundle routing updates.

Default If you omit this statement, routes are exported to BGP immediately after they have been added to the routing table.

Options *seconds*—Output delay time.

Range: 0 through 65,535 seconds

Default: 0 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Applying Policies to BGP Routes

outbound-route-filter

Syntax

```
outbound-route-filter {
  bgp-orf-cisco-mode;
  prefix-based {
    accept {
      (inet | inet6);
    }
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp group *group-name*],
[edit logical-systems *logical-system-name* protocols bgp group *group-name* neighbor *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address*],
[edit protocols bgp],
[edit protocols bgp group *group-name*],
[edit protocols bgp group *group-name* neighbor *address*],
[edit routing-instances *routing-instance-name* protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address*]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure a BGP peer to accept outbound route filters from a remote peer.

Options

accept—Specify that outbound route filters from a BGP peer be accepted.

inet—Specify that IPv4 prefix-based outbound route filters be accepted.

inet6—Specify that IPv6 prefix-based outbound route filters be accepted.

NOTE: You can specify that both IPv4 and IPv6 outbound route filters be accepted.

prefix-based—Specify that prefix-based filters be accepted.

The **bgp-orf-cisco-mode** statement is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Applying Filters Provided by BGP Peers to Outbound Routes

overload

Syntax

```
overload {
  advertise-high-metrics;
  timeout seconds;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols **isis**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **isis**],
[edit protocols **isis**],
[edit routing-instances *routing-instance-name* protocols **isis**]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure the local routing device so that it appears to be overloaded. You might want to do this when you want the routing device to participate in IS-IS routing, but do not want it to be used for transit traffic. Note that traffic to immediately attached interfaces continues to transit the routing device. You can also advertise maximum link metrics in network layer reachability information (NLRI) instead of setting the overload bit.

NOTE: If the time elapsed after the IS-IS instance is enabled is less than the specified timeout, overload mode is set.

Options

advertise-high-metrics—Advertise maximum link metrics in NLRIs instead of setting the overload bit.

timeout *seconds*—Number of seconds at which the overloading is reset.

Default: 0 seconds

Range: 60 through 1800 seconds

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring IS-IS to Make Routing Devices Appear Overloaded

overload

Syntax

```
overload {
    timeout seconds;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols (**ospf** | **ospf3**)],
[edit logical-systems *logical-system-name* protocols ospf topology (default | ipv4-multicast | *name*)],
[edit logical-systems *logical-system-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)],
[edit logical systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf topology (default | ipv4-multicast | *name*)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit protocols (**ospf** | **ospf3**)],
[edit protocols ospf topology (default | ipv4-multicast | *name*)],
[edit protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)]
[edit routing-instances *routing-instance-name* protocols ospf topology (default | ipv4-multicast | *name*],
[edit routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure the local routing device so that it appears to be overloaded. You might do this when you want the routing device to participate in OSPF routing, but do not want it to be used for transit traffic.

NOTE: Traffic destined to immediately attached interfaces continues to reach the routing device.

Options

timeout *seconds*—(Optional) Number of seconds at which the overloading is reset. If no timeout interval is specified, the routing device remains in overload state until the **overload** statement is deleted or a timeout is set.

Range: 60 through 1800 seconds

Default: 0 seconds

NOTE: Multitopology Routing does not support the timeout option.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring OSPF to Make Routing Devices Appear Overloaded
- Configuring Multitopology Routing in OSPF

passive

Syntax
passive;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name* **neighbor** *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],
[edit protocols bgp],
[edit protocols bgp **group** *group-name*],
[edit protocols bgp group *group-name* **neighbor** *address*],
[edit routing-instances *routing-instance-name* protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*]

Release Information
Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Do not send active open messages to the peer. Rather, wait for the peer to issue an open request.

Default
If you omit this statement, all explicitly configured peers are active, and each peer periodically sends open requests until its peer responds.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Disabling Transmission of Open Requests to BGP Peers

passive

Syntax passive;

Hierarchy Level [edit logical-systems *logical-system-name* protocols isis **interface** *interface-name*],
[edit logical-systems *logical-system-name* protocols isis interface *interface-name* level *level-number*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis interface *interface-name* level *level-number*],
[edit protocols isis **interface** *interface-name*],
[edit protocols isis interface *interface-name* level *level-number*],
[edit routing-instances *routing-instance-name* protocols isis **interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols isis interface *interface-name* level *level-number*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Advertise the direct interface addresses on an interface or into a level on the interface without actually running IS-IS on that interface or level.

This statement effectively prevents IS-IS from running on the interface. To enable IS-IS on an interface, include the **interface** statement at the **[edit protocols isis]** or the **[edit routing-instances *routing-instance-name* protocols isis]** hierarchy level. To disable it, include the **disable** statement at those hierarchy levels. The three states are mutually exclusive.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- disable
- Configuring Levels on IS-IS Interfaces

passive

Syntax

```
passive {
    traffic-engineering {
        remote-node-id address;
    }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols (ospf | ospf3) area *area-id* **interface** *interface-name*],
[edit logical-systems *logical-system-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **interface** *interface-name*],
[edit protocols (ospf | ospf3) area *area-id* **interface** *interface-name*],
[edit protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* **interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **interface** *interface-name*]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Advertise the direct interface addresses on an interface without actually running OSPF on that interface. A passive interface is one for which the address information is advertised as an internal route in OSPF, but on which the protocol does not run.

To configure an interface in OSPF passive traffic engineering mode, include the **traffic-engineering** statement. Configuring OSPF passive traffic engineering mode enables the dynamic discovery of OSPF AS boundary routers.

Enable OSPF on an interface by including the **interface** statement at the **[edit protocols (ospf | ospf3) area *area-id*]** or the **[edit routing-instances *routing-instance-name* protocols ospf area *area-id*]** hierarchy levels. Disable it by including the **disable** statement, To prevent OSPF from running on an interface, include the **passive** statement. These three states are mutually exclusive.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **disable on page 1524**
- Advertising Interface Addresses Without Running OSPF
- Configuring OSPF Passive Traffic Engineering Mode

peer-as

Syntax peer-as *autonomous-system*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name* neighbor *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],
[edit protocols bgp],
[edit protocols bgp **group** *group-name*],
[edit protocols bgp group *group-name* **neighbor** *address*],
[edit routing-instances *routing-instance-name* protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the neighbor (peer) AS number.

The autonomous system (AS) numeric range in plain-number format provides BGP support for 4-byte AS numbers, as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*.

You can also configure a 4-byte AS number using the AS-dot notation format of two integer values joined by a period: *<16-bit high-order value in decimal>.<16-bit low-order value in decimal>*. For example, the 4-byte AS number of 65,546 in plain-number format is represented as 1.10 in the AS-dot notation format.

Options ***autonomous-system***—AS number.

Range: 1 through 4,294,967,295 ($2^{32} - 1$) in plain-number format

Range: 0.0 through 65535.65535 in AS-dot notation format

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring BGP Groups and Peers
- Configuring BGP Groups and Peers

pim-to-igmp-proxy

Syntax

```
pim-to-igmp-proxy {
  upstream-interface [ interface-names ];
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure the rendezvous point (RP) routing device that resides between a customer edge–facing Protocol Independent Multicast (PIM) domain and a core-facing PIM domain to translate PIM join or prune messages into corresponding Internet Group Management Protocol (IGMP) report or leave messages. The routing device then transmits the report or leave messages by proxying them to one or two upstream interfaces that you configure on the RP routing device. Including the **pim-to-igmp-proxy** statement enables you to use IGMP to forward IPv4 multicast traffic across the PIM sparse mode domains.

The remaining statement is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM-to-IGMP Message Translation

pim-to-mld-proxy

Syntax

```
pim-to-mld-proxy {
    upstream-interface [ interface-names ];
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure the rendezvous point (RP) routing device that resides between a customer edge–facing Protocol Independent Multicast (PIM) domain and a core-facing PIM domain to translate PIM join or prune messages into corresponding Multicast Listener Discovery (MLD) report or leave messages. The routing device then transmits the report or leave messages by proxying them to one or two upstream interfaces that you configure on the RP routing device. Including the **pim-to-mld-proxy** statement enables you to use MLD to forward IPv6 multicast traffic across the PIM sparse mode domains.

The remaining statement is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM-to-MLD Message Translation

point-to-point

Syntax

```
point-to-point;
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols isis **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis **interface** *interface-name*],
[edit protocols isis **interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols isis **interface** *interface-name*]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure an IS-IS interface to behave like a point-to-point connection.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Point-to-Point Interfaces for IS-IS

policy

Syntax policy *policy-name*;

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options (**aggregate** | **generate**) (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* (**aggregate** | **generate**) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options (**aggregate** | **generate**) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* (**aggregate** | **generate**) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options (**aggregate** | **generate**) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* (**aggregate** | **generate**) (defaults | route)],
[edit routing-options (**aggregate** | **generate**) (defaults | route)],
[edit routing-options rib *routing-table-name* (**aggregate** | **generate**) (defaults | route)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Associate a routing policy when configuring an aggregate or generated route's destination prefix in the **routes** part of the **aggregate** or **generate** statement. This provides the equivalent of an import routing policy filter for the destination prefix. That is, each potential contributor to an aggregate route, along with any aggregate options, is passed through the policy filter. The policy then can accept or reject the route as a contributor to the aggregate route and, if the contributor is accepted, the policy can modify the default preferences. The contributor with the numerically smallest prefix becomes the most preferred, or *primary*, contributor. A rejected contributor still can contribute to a less specific aggregate route. If you do not specify a policy filter, all candidate routes contribute to an aggregate route.

Options ***policy-name***—Name of a routing policy.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **aggregate on page 1470**
- **generate on page 1543**
- Configuring Aggregate Routes
- Configuring Generated Routes

policy (Flow Maps)

Syntax policy [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast **flow-map** *flow-map-name*],
[edit logical-systems *logical-system-name* routing-options multicast **flow-map** *flow-map-name*],
[edit routing-instances *routing-instance-name* routing-options multicast **flow-map** *flow-map-name*],
[edit routing-options multicast **flow-map** *flow-map-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a flow map policy.

Options ***policy-names***—Name of one or more policies for flow mapping.

Required Privilege Level routing—To view this statement in the configuration.

policy (SSM Maps)

Syntax policy [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast **ssm-map** *ssm-map-name*],
[edit logical-systems *logical-system-name* routing-options multicast **ssm-map** *ssm-map-name*],
[edit routing-instances *routing-instance-name* routing-options multicast **ssm-map** *ssm-map-name*],
[edit routing-options multicast **ssm-map** *ssm-map-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply one or more policies to an SSM map.

Options ***policy-names***—Name of one or more policies for SSM mapping.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To view this statement in the configuration.

Related Documentation

- Example: Configuring SSM Mapping

ppm

Syntax

```
ppm {
  centralized;
}
```

Hierarchy Level [edit protocols lacp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure PPM processing options for Link Aggregation Control Protocol (LACP) packets.

This command configures the PPM processing options for LACP packets only. You can disable distributed PPM processing for all packets that use PPM and run all PPM processing on the Routing Engine by entering the **no-delegate-processing** configuration statement in the **[edit routing-options ppm]** statement hierarchy.

Default Distributed PPM processing is enabled for all packets that use PPM.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Distributed Periodic Packet Management on a J-EX Series Switch (CLI Procedure) on page 1451

ppm

Syntax

```
ppm {
  no-delegate-processing;
}
```

Hierarchy Level [edit logical-systems *logical-system-name* routing-options],
[edit routing-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.
no-delegate-processing statement introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Disable distributed periodic packet management (PPM) to the Packet Forwarding Engine (on routers), to access ports (on J-EX4200 switches), or line cards (on J-EX8200 switches).

After you disable PPM, PPM processing continues to run on the Routing Engine.

Default enabled

Options **no-delegate-processing**—Disable PPM to the Packet Forwarding Engine, access ports, or line cards. Distributed PPM is enabled by default.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Disabling Distributed Periodic Packet Management on the Packet Forwarding Engine

preempt

Syntax

```
(preempt | no-preempt) {
    hold-time seconds;
}
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* **vrrp-inet6-group** *group-id*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure whether a backup router can preempt a master router:

- **preempt**—Allow the master router to be preempted.
- **no-preempt**—Prohibit the preemption of the master router.

The remaining statement is explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring VRRP for IPv6 (CLI Procedure) on page 1452

preference

Syntax preference *preference*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name* neighbor *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],
[edit protocols bgp],
[edit protocols bgp **group** *group-name*],
[edit protocols bgp group *group-name* **neighbor** *address*],
[edit routing-instances *routing-instance-name* protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the preference for routes learned from BGP.

At the BGP global level, the preference statement sets the preference for routes learned from BGP. You can override this preference in a BGP group or peer preference statement.

At the group or peer level, the preference statement sets the preference for routes learned from the group or peer. Use this statement to override the preference set in the BGP global preference statement when you want to favor routes from one group or peer over those of another.

Options ***preference***—Preference to assign to routes learned from BGP or from the group or peer.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

Default: 170 for the primary preference

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **local-preference on page 1599**
- Configuring the Default Preference Value for BGP Routes

preference

Syntax preference *preference*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols isis **level** *level-number*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis **level** *level-number*],
[edit protocols isis **level** *level-number*],
[edit routing-instances *routing-instance-name* protocols isis **level** *level-number*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the preference of internal routes.

Options ***preference***—Preference value.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

Default: 15 (for Level 1 internal routes), 18 (for Level 2 internal routes), 160 (for Level 1 external routes), 165 (for Level 2 external routes)

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **external-preference on page 1534**
- Configuring Preference Values for IS-IS Routes

preference

Syntax preference *preference*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols (**ospf** | **ospf3**)],
[edit logical-systems *logical-system-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit protocols (**ospf** | **ospf3**)],
[edit protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)],
[edit routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set the route preference for OSPF internal routes.

Options *preference*—Preference value.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

Default: 10

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **external-preference on page 1535**
- Configuring Preference Values for OSPF Routes

preference

Syntax preference *preference*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols rip **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip **group** *group-name*],
[edit protocols rip **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols rip **group** *group-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the preference of external routes learned by RIP as compared to those learned from other routing protocols.

Options ***preference***—Preference value. A lower value indicates a more preferred route.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

Default: 100

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Group-Specific RIP Properties

preference

Syntax preference *preference*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols ripng **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng group *group-name*],
[edit protocols ripng **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols ripng group *group-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the preference of external routes learned by RIPng as compared to those learned from other routing protocols.

Options ***preference***—Preference value. A lower value indicates a more preferred route.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

Default: 100

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Group-Specific RIPng Properties

preference

Syntax (preference | preference2 | color | color2) *preference* <type *type*>;

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options (**aggregate** | **generate** | **static**) (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* (**aggregate** | **generate** | **static**) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options (**aggregate** | **generate** | **static**) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* (**aggregate** | **generate** | **static**) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options (**aggregate** | **generate** | **static**) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* (**aggregate** | **generate** | **static**) (defaults | route)],
[edit routing-options (**aggregate** | **generate** | **static**) (defaults | route)],
[edit routing-options rib *routing-table-name* (**aggregate** | **generate** | **static**) (defaults | route)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Preference value for a static, aggregated, or generated route. You also can specify a secondary preference value (**preference2**), as well as colors, which are even finer-grained preference values (**color** and **color2**).

Options *preference*—Preference value. A lower number indicates a more preferred route.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

Default: 5 (for static routes), 130 (for aggregate and generated routes)

type—(Optional) Type of route.

Range: 1 through 16

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **aggregate on page 1470**
- **generate on page 1543**
- **static on page 1708**
- Configuring Static Routes
- Configuring Aggregate Routes
- Configuring Generated Routes

prefix

Syntax prefix *destination-prefix*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast scope *scope-name*],
[edit logical-systems *logical-system-name* routing-options multicast scope *scope-name*],
[edit routing-instances *routing-instance-name* routing-options multicast scope *scope-name*],
[edit routing-options multicast scope *scope-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the prefix for multicast scopes.

Options ***destination-prefix***—Address range for the multicast scope.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **multicast on page 1621**
- Configuring Multicast Scoping

prefix-export-limit

Syntax prefix-export-limit *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols isis **level** *level-number*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis **level** *level-number*],
[edit protocols isis **level** *level-number*],
[edit routing-instances *routing-instance-name* protocols isis **level** *level-number*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a limit to the number of prefixes exported into IS-IS.

Options ***number***—Prefix limit.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Limiting the Number of Prefixes Exported to IS-IS

prefix-export-limit

Syntax prefix-export-limit *number*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols (**ospf** | **ospf3**)],
[edit logical-systems *logical-system-name* protocols ospf topology (default | ipv4-multicast | *name*)],
[edit logical-systems *logical-system-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf topology (default | ipv4-multicast | *name*)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit protocols (**ospf** | **ospf3**)],
[edit protocols ospf topology (default | ipv4-multicast | *name*)],
[edit protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)],
[edit routing-instances *routing-instance-name* protocols ospf topology (default | ipv4-multicast | *name*)],
[edit routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a limit to the number of prefixes exported into OSPF.

Options *number*—Prefix limit.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

Default: None

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Limiting the Number of Prefixes Exported to OSPF
- Configuring Multitopology Routing in OSPF

prefix-limit

Syntax

```
prefix-limit {
    maximum number;
    teardown <percentage> <idle-timeout (forever | minutes)>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols bgp **family** (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)],
[edit logical-systems *logical-system-name* protocols bgp group *group-name* **family** (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)],
[edit logical-systems *logical-system-name* protocols bgp group *group-name* neighbor *address* **family** (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **family** (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **family** (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address* **family** (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)],
[edit protocols bgp **family** (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)],
[edit protocols bgp group *group-name* **family** (inet | inet6) (any | labeled-unicast | multicast | unicast)],
[edit protocols bgp group *group-name* neighbor *address* **family** (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)],
[edit routing-instances *routing-instance-name* protocols bgp **family** (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **family** (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address* **family** (inet | inet6) (any | flow | labeled-unicast | multicast | unicast)]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Limit the number of prefixes received on a BGP peering session and a rate-limit logging when injected prefixes exceed a set limit.

Options

maximum *number*—When you set the maximum number of prefixes, a message is logged when that number is exceeded.

Range: 1 through 4,294,967,295 ($2^{32} - 1$)

teardown <*percentage*>—If you include the **teardown** statement, the session is torn down when the maximum number of prefixes is reached. If you specify a percentage, messages are logged when the number of prefixes exceeds that percentage. After the session is torn down, it is reestablished in a short time unless you include the **idle-timeout** statement. Then the session can be kept down for a specified amount of time, or forever. If you specify **forever**, the session is reestablished only after you issue a **clear bgp neighbor** command.

Range: 1 through 100

idle-timeout (forever | *timeout-in-minutes*)—(Optional) If you include the **idle-timeout** statement, the session is torn down for a specified amount of time, or forever. If you specify a period of time, the session is allowed to reestablish after this timeout period. If you specify **forever**, the session is reestablished only after you intervene with a **clear bgp neighbor** command.

Range: 1 through 2400

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- accepted-prefix-limit
- Enabling Multiprotocol BGP

priority

Syntax
priority *number*;

Hierarchy Level
[edit interfaces *interface-name* unit *logical-unit-number* family inet address *address* vrrp-group *group-id*],
[edit interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* **vrrp-inet6-group** *group-id*]

Release Information
Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Configure a switch's priority for becoming the master default routing platform. The routing platform with the highest priority within the group becomes the master.

Options
number—Routing platform's priority for being elected to be the master router in the VRRP group. A larger value indicates a higher priority for being elected.

Range: 1 through 255

Default: 100 (for backup routers)

NOTE: Priority 255 cannot be assigned to routed VLAN interfaces (RVIs).

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring VRRP for IPv6 (CLI Procedure) on page 1452

priority

Syntax priority *number*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols isis interface *interface-name* level *level-number*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis interface *interface-name* level *level-number*],
[edit protocols isis interface *interface-name* level *level-number*],
[edit routing-instances *routing-instance-name* protocols isis interface *interface-name* level *level-number*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description The interface's priority for becoming the designated router. The interface with the highest priority value becomes that level's designated router.

The priority value is meaningful only on a multiaccess network. It has no meaning on a point-to-point interface.

Options *number*—Priority value.

Range: 0 through 127

Default: 64

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Levels on IS-IS Interfaces

priority

Syntax priority *number*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols (ospf | ospf3) area *area-id* interface *interface-name*],
[edit logical-systems *logical-system-name* protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)] area *area-id* interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* interface *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)] area *area-id* interface *interface-name*],
[edit protocols (ospf | ospf3) area *area-id* interface *interface-name*],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)] area *area-id* interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* interface *interface-name*],
[edit routing-instances *routing-instance-name* protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)] area *area-id* interface *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the routing device's priority for becoming the designated routing devices. The routing device that has the highest priority value on the logical IP network or subnet becomes the network's designated router. You must configure at least one routing device on each logical IP network or subnet to be the designated router. You also should specify a routing device's priority for becoming the designated router on point-to-point interfaces.

Options *number*—Routing device's priority for becoming the designated router. A priority value of 0 means that the routing device never becomes the designated router. A value of 1 means that the routing device has the least chance of becoming a designated router.

Range: 0 through 255

Default: 128

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- OSPF Designated Router Overview
- Configuring the Designated Router Priority for OSPF

qualified-next-hop

Syntax

```
qualified-next-hop (address | interface-name) {
    interface interface-name;
    metric metric;
    preference preference;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* **routing-options static** route *destination-prefix*],
[edit logical-systems *logical-system-name* **routing-options** rib inet6.0 **static** route *destination-prefix*],
[edit logical-systems *logical-system-name* **routing-options static** route *destination-prefix*],
[edit routing-instances *routing-instance-name* **routing-options static** route *destination-prefix*],
[edit **routing-options** rib inet6.0 **static** route *destination-prefix*],
[edit **routing-options static** route *destination-prefix*]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure an independent metric or preference on a static route.

Options

address—IPv4, IPv6, or ISO network address of the next hop.

interface-name—Name of the interface on which to configure an independent metric or preference for a static route. To configure an unnumbered Ethernet interface as the next-hop interface for a static route, specify **qualified-next-hop *interface-name***, where ***interface-name*** is the name of the IPv4 or IPv6 unnumbered Ethernet interface.

metric—Metric value.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

preference—Preference value. A lower number indicates a more preferred route.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

Default: 5

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring an Independent Preference for Static Routes

readvertise

Syntax (readvertise | no-readvertise);

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* **static** (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options **static** (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* **static** (defaults | route)],
[edit logical-systems *logical-system-name* routing-options **static** (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* **static** (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options **static** (defaults | route)],
[edit routing-options rib *routing-table-name* **static** (defaults | route)],
[edit routing-options **static** (defaults | route)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure whether static routes are eligible to be readvertised by routing protocols:

- **readvertise**—Readvertise static routes.
- **no-readvertise**—Mark a static route as being ineligible for readvertisement; include the **no-readvertise** option when configuring the route.

Default readvertise

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **static on page 1708**
- Configuring Static Routes

realm

Syntax

```
realm (ipv4-unicast | ipv4-multicast | ipv6-unicast) {
    area area-id {
        interface interface-name;
    }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols **ospf3**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **ospf3**],
[edit protocols **ospf3**],
[edit routing-instances *routing-instance-name* protocols **ospf3**]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure OSPFv3 to advertise address families other than unicast IPv6. The Junos OS maps each address family you configure to a separate realm with its own set of neighbors and link-state database.

Options

ipv4-unicast—Configure a realm for IPv4 unicast routes.

ipv4-multicast—Configure a realm for IPv4 multicast routes.

ipv6-multicast—Configure a realm for IPv6 multicast routes.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Multiple Address Families for OSPFv3

receive

Syntax receive *receive-options*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols **rip**],
[edit logical-systems *logical-system-name* protocols rip group *group-name* **neighbor** *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **rip**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip group *group-name* **neighbor** *neighbor-name*],
[edit protocols **rip**],
[edit protocols rip group *group-name* **neighbor** *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols **rip**],
[edit routing-instances *routing-instance-name* protocols rip group *group-name* **neighbor** *neighbor-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure RIP receive options.

Options ***receive-options***—One of the following:

- **both**—Accept both RIP version 1 and version 2 packets.
- **none**—Do not receive RIP packets.
- **version-1**—Accept only RIP version 1 packets.
- **version-2**—Accept only RIP version 2 packets.

Default: both

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **send on page 1700**
- Configuring RIP Update Messages

receive

Syntax receive <none>;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols ripng],
[edit logical-systems *logical-system-name* protocols ripng group *group-name* neighbor *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng group *group-name* neighbor *neighbor-name*],
[edit protocols ripng],
[edit protocols ripng group *group-name* neighbor *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols ripng],
[edit routing-instances *routing-instance-name* protocols ripng group *group-name* neighbor *neighbor-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable or disable receiving of update messages.

Options **none**—(Optional) Disable receiving update messages.

Default: Enabled

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **send on page 1701**
- Configuring RIPng Update Messages

redundant-sources

Syntax redundant-sources [*addresses*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast **flow-map** *flow-map-name*],
[edit logical-systems *logical-system-name* routing-options multicast **flow-map** *flow-map-name*],
[edit routing-instances *routing-instance-name* routing-options multicast **flow-map** *flow-map-name*],
[edit routing-options multicast **flow-map** *flow-map-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a list of redundant sources for multicast flows defined by a flow map.

Options ***addresses***—List of IPv4 or IPv6 addresses for use as redundant (backup) sources for multicast flows defined by a flow map.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Defining Flow Properties

reference-bandwidth

Syntax reference-bandwidth *reference-bandwidth*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols **isis**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **isis**],
[edit protocols **isis**],
[edit routing-instances *routing-instance-name* protocols **isis**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set the reference bandwidth used in calculating the default interface cost. The cost is calculated using the following formula:

```
cost = reference-bandwidth/bandwidth
```

Options ***reference-bandwidth***—Reference bandwidth, in megabits per second.

Default: 10 Mbps

Range: 9600 through 1,000,000,000,000 Mbps

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring the Reference Bandwidth Used in IS-IS Metric Calculations

reference-bandwidth

Syntax reference-bandwidth *reference-bandwidth*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols (**ospf** | **ospf3**)],
[edit logical-systems *logical-system-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)],
[edit logical-systems *logical-system-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit protocols (**ospf** | **ospf3**)],
[edit protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)],
[edit routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set the reference bandwidth used in calculating the default interface cost. The cost is calculated using the following formula:

```
cost = ref-bandwidth/bandwidth
```

Options ***ref-bandwidth***—Reference bandwidth, in bits per second.

Default: 100 Mbps (100,000,000 bits)

Range: 9600 through 1,000,000,000,000 bits

NOTE: The default behavior is to use the reference-bandwidth value to calculate the cost of OSPF interfaces. You can override this behavior for any OSPF interface by configuring a specific cost with the metric statement.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **metric on page 1611**
- Configuring the Metric Value for OSPF Interfaces

remove-private

Syntax remove-private;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name* neighbor *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],
[edit protocols bgp],
[edit protocols bgp **group** *group-name*],
[edit protocols bgp group *group-name* **neighbor** *address*],
[edit routing-instances *routing-instance-name* protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description When advertising AS paths to remote systems, have the local system strip private AS numbers from the AS path. The numbers are stripped from the AS path starting at the left end of the AS path (the end where AS paths have been most recently added). The routing device stops searching for private ASs when it finds the first nonprivate AS or a peer's private AS. This operation takes place after any confederation member ASs have already been removed from the AS path, if applicable.

The Junos OS recognizes the set of AS numbers that is considered private, a range that is defined in the Internet Assigned Numbers Authority (IANA) assigned numbers document.

The set of reserved AS numbers is in the range from 64,512 through 65,535.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Removing Private AS Numbers from AS Paths

resolution

Syntax

```
resolution {
  rib routing-table-name {
    import [ policy-names ];
    resolution-ribs [ routing-table-names ];
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure route resolution.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Route Resolution

resolution-ribs

Syntax

```
resolution-ribs [ routing-table-names ];
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options resolution rib],
[edit logical-systems *logical-system-name* routing-options resolution rib],
[edit routing-instances *routing-instance-name* routing-options resolution rib],
[edit routing-options resolution rib]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify one or more routing tables to use for route resolution.

The remaining statements are explained separately.

Options ***routing-table-names***—Name of one or more routing tables.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Route Resolution

resolve

Syntax resolve;

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* **static** (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options **static** (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* **static** (defaults | route)],
[edit logical-systems *logical-system-name* routing-options **static** (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* **static** (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options **static** (defaults | route)],
[edit routing-options rib *routing-table-name* **static** (defaults | route)],
[edit routing-options **static** (defaults | route)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure statically configured routes to be resolved to a next hop that is not directly connected. The route is resolved through the **inet.0** and **inet.3** routing tables.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **static on page 1708**
- Configuring Static Route Options

restart-duration

Syntax restart-duration *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* **routing-options** graceful-restart],
[edit logical-systems *logical-system-name* **routing-options** graceful-restart],
[edit routing-instances *routing-instance-name* **routing-options** graceful-restart],
[edit **routing-options** graceful-restart]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the restart timer for graceful restart.

Options **restart-duration *seconds***—Configure the time period for the restart to last.

Range: 120 through 900 seconds

Default: 90 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Graceful Restart

retain

Syntax (retain | no-retain);

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* **static** (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options **static** (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* **static** (defaults | route)],
[edit logical-systems *logical-system-name* routing-options **static** (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* **static** (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options **static** (defaults | route)],
[edit routing-options rib *routing-table-name* **static** (defaults | route)],
[edit routing-options **static** (defaults | route)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure statically configured routes to be deleted from or retained in the forwarding table when the routing protocol process shuts down normally:

- **retain**—Have a static route remain in the forwarding table when the routing protocol process shuts down normally. Doing this greatly reduces the time required to restart a system that has a large number of routes in its routing table.
- **no-retain**—Delete statically configured routes from the forwarding table when the routing protocol process shuts down normally.

Default **no-retain**

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **static on page 1708**
- Configuring Static Routes

retransmit-interval

Syntax retransmit-interval *seconds*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols ospf area *area-id* peer-interface *interface-name*],
[edit logical-systems *logical-system-name* protocols (ospf | ospf3) area *area-id* **interface** *interface-name*],
[edit logical-systems *logical-system-name* protocols (ospf | ospf3) area *area-id* **virtual-link**],
[edit logical-systems *logical-system-name* protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* **virtual-link**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **interface** *interface-name*],
[edit protocols ospf area *area-id* peer-interface *interface-name*],
[edit protocols (ospf | ospf3) area *area-id* **interface** *interface-name*],
[edit protocols (ospf | ospf3) area *area-id* **virtual-link**],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* **interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* **virtual-link**],
[edit routing-instances *routing-instance-name* protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify how long the routing device waits to receive a link-state acknowledgment packet before retransmitting link-state advertisements to an interface's neighbors.

Options ***seconds***—Interval to wait.

Range: 1 through 65,535 seconds

Default: 5 seconds

NOTE: You must configure link-state advertisement (LSA) retransmit intervals to be equal to or greater than 3 seconds to avoid triggering a retransmit trap, because the Junos OS delays LSA acknowledgments by up to 2 seconds.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring OSPF Timers

reverse-oif-mapping

Syntax

```
reverse-oif-mapping {
    no-qos-adjust;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-options multicast **interface** *interface-name*],
[edit routing-instances *routing-instance-name* routing-options multicast **interface** *interface-name*],
[edit routing-options multicast **interface** *interface-name*]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Enable the routing device to identify a subscriber VLAN or interface based on an IGMP or MLD request it receives over the multicast VLAN.

The remaining statement is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Managing Subscriber Overcommitment

rib (General)

Syntax

```
rib routing-table-name {
  aggregate {
    defaults {
      ... aggregate-options ...
    }
    route destination-prefix {
      policy policy-name;
      ... aggregate-options ...
    }
    generate {
      defaults {
        generate-options;
      }
      route destination-prefix {
        policy policy-name;
        generate-options;
      }
    }
    martians {
      destination-prefix match-type <allow>;
    }
  }
  static {
    defaults {
      static-options;
    }
    rib-group group-name;
    route destination-prefix {
      next-hop;
      static-options;
    }
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Create a routing table.

Explicitly creating a routing table with the ***routing-table-name*** statement is optional if you are not adding any static, martian, aggregate, or generated routes to the routing table and if you also are creating a routing table group. Simply including the **passive** statement to indicate that a routing table is part of a routing table group is sufficient to create it.

NOTE: The IPv4 multicast routing table (inet.1) and the IPv6 multicast routing table (inet6.1) are not supported for this statement.

Default If you do not specify a routing table name with the ***routing-table-name*** statement, the software uses the default routing tables, which are **inet.0** for unicast routes and **inet.1** for the multicast cache.

Options ***routing-table-name***—Name of the routing table, in the following format:
protocol [***.identifier***]

- ***protocol*** is the protocol family. It can be **inet6** for the IPv6 family, **inet** for the IPv4 family, **iso** for the ISO protocol family, or ***instance-name*.iso.0** for an ISO routing instance.
- ***identifier*** is a positive integer that specifies the instance of the routing table.

Default: **inet.0**

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- passive
- Creating Routing Tables

rib (Route Resolution)

Syntax
```
rib routing-table-name {
  import [ policy-names ];
  resolution-ribs [ routing-table-names ];
}
```

Hierarchy Level
```
[edit logical-systems logical-system-name routing-instances routing-instance-name
  routing-options resolution],
[edit logical-systems logical-system-name routing-options resolution],
[edit routing-instances routing-instance-name routing-options resolution],
[edit routing-options resolution]
```

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify a routing table name for route resolution.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Route Resolution

rib-group

Syntax rib-group *group-name*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols bgp **family** inet (any | labeled-unicast | unicast | multicast)],
[edit logical-systems *logical-system-name* protocols bgp group *group-name* **family** inet (any | labeled-unicast | unicast | multicast)],
[edit logical-systems *logical-system-name* protocols bgp group *group-name* neighbor *address* **family** inet (any | labeled-unicast | unicast | multicast)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **family** inet (any | labeled-unicast | unicast | multicast)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **family** inet (any | labeled-unicast | unicast | multicast)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address* **family** inet (any | labeled-unicast | unicast | multicast)],
[edit protocols bgp **family** inet (any | labeled-unicast | unicast | multicast)],
[edit protocols bgp group *group-name* **family** inet (any | labeled-unicast | unicast | multicast)],
[edit protocols bgp group *group-name* neighbor *address* **family** inet (any | labeled-unicast | unicast | multicast)],
[edit routing-instances *routing-instance-name* protocols bgp **family** inet (any | labeled-unicast | unicast | multicast)],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **family** inet (any | labeled-unicast | unicast | multicast)],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address* **family** inet (any | labeled-unicast | unicast | multicast)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Add unicast prefixes to unicast and multicast tables.

Options
group-name—Name of the routing table group. The name must start with a letter and can include letters, numbers, and hyphens. You generally specify only one routing table group.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **interface-routes on page 1585**
- **rib-group on page 1691**
- Creating Routing Table Groups
- Configuring How Interface Routes Are Imported into Routing Tables
- Enabling Multiprotocol BGP

rib-group

Syntax

```
rib-group {
    inet group-name;
    inet6 group-name;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols **isis**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **isis**],
[edit protocols **isis**],
[edit routing-instances *routing-instance-name* protocols **isis**]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Install routes learned from IS-IS routing instances into routing tables in the IS-IS routing table group. You can install IPv4 routes or IPv6 routes.

Support for IPv6 routing table groups in IS-IS enables IPv6 routes that are learned from IS-IS routing instances to be installed into other routing tables defined in an IS-IS routing table group.

Options

group-name—Name of the routing table group.

inet—Install IPv4 IS-IS routes.

inet6—Install IPv6 IS-IS routes.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Creating Routing Table Groups
- Configuring How Interface Routes Are Imported into Routing Tables
- Enabling Multiprotocol BGP

rib-group

Syntax rib-group *group-name*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols (**ospf** | **ospf3**)],
[edit logical-systems *logical-system-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit protocols (**ospf** | **ospf3**)],
[edit protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)],
[edit routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Install routes learned from OSPF routing instances into routing tables in the OSPF routing table group.

Options *group-name*—Name of the routing table group.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **interface-routes on page 1585**
- **rib-group on page 1691**
- Creating Routing Table Groups
- Configuring How Interface Routes Are Imported into Routing Tables
- Enabling Multiprotocol BGP

rib-group

Syntax rib-group *group-name*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols **rip**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **rip**],
[edit protocols **rip**],
[edit routing-instances *routing-instance-name* protocols **rip**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Install RIP routes into multiple routing tables by configuring a routing table group.

Options ***group-name***—Name of the routing table group.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Routing Table Groups for RIP

rib-group

Syntax rib-group *group-name*;

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options **interface-routes**],
[edit logical-systems *logical-system-name* routing-options **interface-routes**],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* static],
[edit logical-systems *logical-system-name* routing-options static],
[edit routing-instances *routing-instance-name* routing-options **interface-routes**],
[edit routing-options **interface-routes**],
[edit routing-options rib *routing-table-name* static],
[edit routing-options static]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure which routing table groups interface routes are imported into.

Options ***group-name***—Name of the routing table group. The name must start with a letter and can include letters, numbers, and hyphens. It generally does not make sense to specify more than a single routing table group.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **interface-routes on page 1585**
- **rib-groups on page 1692**
- Configuring How Interface Routes Are Imported into Routing Tables
- Creating Routing Table Groups

rib-groups

Syntax

```
rib-groups {
  group-name {
    export-rib group-name;
    import-policy [ policy-names ];
    import-rib [ group-names ];
  }
}
```

Hierarchy Level [edit logical-systems *logical-system-name* routing-options],
[edit routing-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Group one or more routing tables to form a routing table group. A routing protocol can import routes into all the routing tables in the group and can export routes from a single routing table.

Each routing table group must contain one or more routing tables that the Junos OS uses when importing routes (specified in the **import-rib** statement) and optionally can contain one routing table group that the Junos OS uses when exporting routes to the routing protocols (specified in the **export-rib** statement).

Options ***group-name***—Name of the routing table group. The name must start with a letter and can include letters, numbers, and hyphens.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **rib-group on page 1691**
- Creating Routing Table Groups

rip

Syntax rip {...}

Hierarchy Level [edit logical-systems *logical-system-name* protocols],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols],
[edit protocols],
[edit routing-instances *routing-instance-name* protocols]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable RIP routing on the routing device.

Default RIP is disabled on the routing device.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Minimum RIP Configuration

ripng

Syntax ripng {...}

Hierarchy Level [edit logical-systems *logical-system-name* protocols],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols],
[edit protocols],
[edit routing-instances *routing-instance-name* protocols]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable RIPng routing on the routing device.

Default RIPng is disabled on the routing device.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Minimum RIPng Configuration

route-distinguisher-id

Syntax route-distinguisher-id *address*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a route distinguisher identifier for a routing instance, specifying an IP address. If a route distinguisher is configured for a particular routing instance, that value supersedes the route distinguisher configured by this statement.

Options ***address***—IP address.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Route Distinguishers for VRF and Layer 2 VPN Instances

route-record

Syntax route-record;

Hierarchy Level [edit logical-systems *logical-system-name* routing-options],
[edit routing-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Export the AS path and routing information to the traffic sampling process.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Route Recording for Flow Aggregation
- *Junos OS Network Interfaces Configuration Guide*

route-timeout

Syntax route-timeout *seconds*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols **rip**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **rip**],
[edit logical-systems *logical-system-name* protocols rip **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip **group** *group-name*],
[edit protocols **rip**],
[edit protocols rip **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols **rip**],
[edit routing-instances *routing-instance-name* protocols rip **group** *group-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the route timeout interval for RIP.

Options ***seconds***—Estimated time to wait before making updates to the routing table.

Range: 30 through 360 seconds

Default: 180 seconds

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring RIP Timers

route-timeout

Syntax route-timeout *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols ripng],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng],
[edit protocols ripng],
[edit routing-instances *routing-instance-name* protocols ripng]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the route timeout interval for RIPng.

Options ***seconds***—Estimated time to wait before making updates to the routing table.

Range: 30 through 360 seconds

Default: 180 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring RIPng Timers

route-type-community

Syntax route-type-community (iana | vendor);

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)],
[edit routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify an extended community value to encode the OSPF route type. Each extended community is coded as an eight-octet value. This statement sets the most significant bit to either an IANA or vendor-specific route type.

Options **iana**—Encode a route type with the value **0x0306**. This is the default value.

vendor—Encode the route type with the value **0x8000**.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring OSPF Domain IDs for VPNs

router-id

Syntax router-id *address*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the routing device's IP address.

NOTE: We strongly recommend that you configure the router identifier under the [edit routing-options] hierarchy level to avoid unpredictable behavior if the interface address on a loopback interface changes.

Options ***address***—IP address of the routing device.

Default: Address of the first interface encountered by the Junos OS

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Router Identifiers for BGP and OSPF

routing-options

Syntax routing-options { ... }

Hierarchy Level [edit],
[edit logical-systems *logical-system-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name*],
[edit routing-instances *routing-instance-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure protocol-independent routing properties.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Protocol-Independent Routing Properties Configuration Statements

rpf-check-policy

Syntax rpf-check-policy [*policy-names*];

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply policies for disabling RPF checks on arriving multicast packets. The policies must be correctly configured.

Options ***policy-names***—Name of one or more multicast RPF check policies.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring RPF Policies

scope

Syntax
```
scope scope-name {
  interface [ interface-names ];
  prefix destination-prefix;
}
```

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure multicast scoping.

Options ***scope-name***—Name of the multicast scope.

The remaining statements are explained separately.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Creating a Named Scope for Multicast Scoping

scope-policy

Syntax scope-policy [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-options multicast]

NOTE: You can configure a scope policy at these two hierarchy levels only. You cannot apply a scope policy to a specific routing instance, because all scoping policies are applied to all routing instances. However, you can apply the scope statement to a specific routing instance at the [edit routing-instances *routing-instance-name* routing-options multicast] or [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast] hierarchy level.

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply policies for scoping. The policy must be correctly configured at the **edit policy-options policy-statement** hierarchy level.

Options ***policy-names***—Name of one or more multicast scope policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **scope on page 1698**
- Using a Scope Policy for Multicast Scoping

send

Syntax send *send-options*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols **rip**],
[edit logical-systems *logical-system-name* protocols rip group *group-name* **neighbor** *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **rip**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip group *group-name* **neighbor** *neighbor-name*],
[edit protocols **rip**],
[edit protocols rip group *group-name* **neighbor** *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols **rip**],
[edit routing-instances *routing-instance-name* protocols rip group *group-name* **neighbor** *neighbor-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure RIP send options.

Options ***send-options***—One of the following:

- **broadcast**—Broadcast RIP version 2 packets (RIP version 1 compatible).
- **multicast**—Multicast RIP version 2 packets. This is the default.
- **none**—Do not send RIP updates.
- **version-1**—Broadcast RIP version 1 packets.

Default: multicast

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **receive on page 1674**
- Configuring RIP Update Messages

send

Syntax send <none>;

Hierarchy Level [edit logical-systems *logical-system-name* protocols ripng],
[edit logical-systems *logical-system-name* protocols ripng group *group-name* neighbor *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* protocols ripng],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng group *group-name* neighbor *neighbor-name*],
[edit protocols ripng],
[edit protocols ripng group *group-name* neighbor *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols ripng],
[edit routing-instances *routing-instance-name* protocols ripng group *group-name* neighbor *neighbor-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable or disable sending of update messages.

Options **none**—(Optional) Disable sending of update messages.

Default: Enabled

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **receive on page 1675**
- Configuring RIPng Update Messages

shortcuts

Syntax

```
shortcuts;
  lsp-metric-into-summary;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols (ospf | ospf3) **traffic-engineering**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf | ospf3) **traffic-engineering**],
[edit protocols (ospf | ospf3) **traffic-engineering**],
[edit routing-instances *routing-instance-name* protocols (ospf | ospf3)**traffic-engineering**]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure OSPF to use MPLS label-switched paths (LSPs) as shortcut next hops. By default, shortcut routes calculated through OSPFv2 are installed in the**inet.3** routing table, and shortcut routes calculated through OSPFv3 are installed in the **inet6.3** routing table.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Enabling OSPF Traffic Engineering Support

source

Syntax

source [*addresses*];

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast **ssm-map** *ssm-map-name*],
[edit logical-systems *logical-system-name* routing-options multicast **ssm-map** *ssm-map-name*],
[edit routing-instances *routing-instance-name* routing-options multicast **ssm-map** *ssm-map-name*],
[edit routing-options multicast **ssm-map** *ssm-map-name*]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Specify IPv4 or IPv6 source addresses for an SSM map.

Options

addresses—IPv4 or IPv6 source addresses.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To view this statement in the configuration.

Related Documentation

- Example: Configuring SSM Mapping

source-routing

Syntax

```
source-routing {
   (ip | ipv6)
}
```

Hierarchy Level [edit routing-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable source routing.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Enabling Source Routing

spf-options

Syntax

```
spf-options {
    delay milliseconds;
    holddown milliseconds;
    rapid-runs number;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols isis],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis],
[edit protocols isis],
[edit routing-instances *routing-instance-name* protocols isis]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure options for running the shortest-path-first (SPF) algorithm. You can configure a delay for when to run the SPF algorithm after a network topology change is detected, the maximum number of times the SPF algorithm can run in succession, and a holddown interval after SPF algorithm runs the maximum number of times.

Options

delay *milliseconds*—Time interval between the detection of a topology change and when the SPF algorithm runs.

Range: 50 through 1000 milliseconds

Default: 200 milliseconds

holddown *milliseconds*—Time interval to hold down, or wait before a subsequent SPF algorithm runs after the SPF algorithm has run the configured maximum number of times in succession.

Range: 2000 through 10,000 milliseconds

Default: 5000 milliseconds

rapid-runs *number*—Maximum number of times the SPF algorithm can run in succession. After the maximum is reached, the holddown interval begins.

Range: 1 through 5

Default: 3

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring SPF Options for IS-IS

spf-options

Syntax

```
spf-options {
    delay milliseconds;
    holddown milliseconds;
    rapid-runs number;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols (**ospf** | **ospf3**)],
[edit logical-systems *logical-system-name* protocols ospf topology (default | ipv4-multicast | *name*)],
[edit logical-systems *logical-system-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf topology (default | ipv4-multicast | *name*)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit protocols (**ospf** | **ospf3**)],
[edit protocols ospf topology (default | ipv4-multicast | *name*)],
[edit protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)],
[edit routing-instances *routing-instance-name* protocols ospf topology (default | ipv4-multicast | *name*)],
[edit routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure options for running the shortest-path-first (SPF) algorithm. You can configure a delay for when to run the SPF algorithm after a network topology change is detected, the maximum number of times the SPF algorithm can run in succession, and a hold-down interval after the SPF algorithm runs the maximum number of times.

Options

delay ***milliseconds***—Time interval between the detection of a topology change and when the SPF algorithm runs.

Range: 50 through 8000 milliseconds

Default: 200 milliseconds

holddown ***milliseconds***—Time interval to hold down, or wait before a subsequent SPF algorithm runs after the SPF algorithm has run the configured maximum number of times in succession.

Range: 2000 through 20,000 milliseconds

Default: 5000 milliseconds

rapid-runs ***number***—Maximum number of times the SPF algorithm can run in succession. After the maximum is reached, the holddown interval begins.

Range: 1 through 5

Default: 3

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring SPF Options for OSPF
- Configuring Multitopology Routing in OSPF

ssm-groups

Syntax
ssm-groups [*ip-addresses*];

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information
Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Configure additional source-specific multicast (SSM) groups.

Options
ip-addresses—List of one or more additional SSM group addresses separated by a space.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Source-Specific Multicast Groups

ssm-map

Syntax

```
ssm-map ssm-map-name {
    policy [ policy-names ];
    source [ addresses ];
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast],
[edit logical-systems *logical-system-name* routing-options multicast],
[edit routing-instances *routing-instance-name* routing-options multicast],
[edit routing-options multicast]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure SSM mapping.

Options

ssm-map-name—Name of the SSM map.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring SSM Mapping

static

Syntax

```
static {
  defaults {
    static-options;
  }
  rib-group group-name;
  route destination-prefix {
    bfd-liveness-detection {
      authentication {
        algorithm algorithm-name;
        key-chain key-chain-name;
        loose-check;
      }
      detection-time {
        threshold milliseconds;
      }
      local-address ip-address;
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      minimum-receive-ttl number;
      multiplier number;
      neighbor address;
      no-adaptation;
      transmit-interval {
        threshold milliseconds;
        minimum-interval milliseconds;
      }
      version (1 | automatic);
    }
    next-hop address;
    next-hop options;
    qualified-next-hop address {
      metric metric;
      preference preference;
    }
    static-options;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name*],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options],
[edit routing-options rib *routing-table-name*]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure static routes to be installed in the routing table. You can specify any number of routes within a single **static** statement, and you can specify any number of **static** options in the configuration.

Options **defaults**—Specify global static route options. These options only set default attributes inherited by all newly created static routes. These are treated as global defaults and apply to all the static routes you configure in the **static** statement. This part of the **static** statement is optional.

route ***destination-prefix***—Destination of the static route.

- **defaults**—For the default route to the destination. This is equivalent to specifying an IP address of **0.0.0.0/0**.
- ***destination-prefix/prefix-length***—***destination-prefix*** is the network portion of the IP address, and ***prefix-length*** is the destination prefix length.
- **next-hop** ***address***—Reach the next-hop routing device by specifying an IP address, an interface name, or an ISO network entity title (NET).
- ***nsap-prefix***—***nsap-prefix*** is the network service access point (NSAP) address for ISO.

next-hop options—Additional information for how to manage forwarding of packets to the next hop.

- **discard**—Do not forward packets addressed to this destination. Instead, drop the packets, do not send ICMP unreachable messages to the packets' originators, and install a reject route for this destination into the routing table.
- ***iso-net***—Reach the next-hop routing device by specifying an ISO NSAP.
- **next-table** ***routing-table-name***—Name of the next routing table to the destination.
- **receive**—Install a receive route for this destination into the routing table.
- **reject**—Do not forward packets addressed to this destination. Instead, drop the packets, send ICMP unreachable messages to the packets' originators, and install a reject route for this destination into the routing table.

static-options—(Optional under **route**) Additional information about static routes, which is included with the route when it is installed in the routing table.

You can specify one or more of the following in ***static-options***. Each of the options is explained separately.

- **(active | passive)**;
- **as-path <***as-path***> <origin (egp | igp | incomplete)> <atomic-aggregate> <aggregator** ***as-number in-address*****>**;
- **community [** ***community-ids*** **]**;
- **(install | no-install)**;
- **(metric | metric2 | metric3 | metric4)** ***value*** **<type** ***type*****>**;
- **(preference | preference2 | color | color2)** ***preference*** **<type** ***type*****>**;
- **(readvertise | no-readvertise)**;

- **(resolve | no-resolve)**;
- **(no-retain | retain)**;
- **tag** *string*;

The remaining statements are explained separately.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Static Routes

stub

Syntax
stub <**default-metric** *metric*> <(no-summaries | summaries)>;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols (ospf | ospf3) **area** *area-id*],
[edit logical-systems *logical-system-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf | ospf3) **area** *area-id*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit protocols (ospf | ospf3) **area** *area-id*],
[edit protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances *routing-instance-name* protocols (ospf | ospf3) **area** *area-id*],
[edit routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)]

Release Information
Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Specify that this area not be flooded with AS external link-state advertisements (LSA)s. You must include the **stub** statement when configuring all routing devices that are in the stub area.

The backbone cannot be configured as a stub area.

You cannot configure an area to be both a stub area and a not-so-stubby area (NSSA).

Options
no-summaries—(Optional) Do not advertise routes into the stub area. If you include the **default-metric** option, only the default route is advertised.

summaries—(Optional) Flood summary LSAs into the stub area.

The remaining statement is explained separately.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **nssa on page 1642**
- Configuring OSPF Areas

subscriber-leave-timer

Syntax subscriber-leave-timer *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-options multicast **interface** *interface-name*],
[edit routing-instances *routing-instance-name* routing-options multicast **interface** *interface-name*],
[edit routing-options multicast **interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Length of time before the multicast VLAN updates QoS data (for example, available bandwidth) for subscriber interfaces after it receives an IGMP leave message.

Options ***seconds***—Length of time before the multicast VLAN updates QoS data (for example, available bandwidth) for subscriber interfaces after it receives an IGMP leave message. Specifying a value of 0 results in an immediate update; this is the same as if the statement were not configured.

Range: 0 through 30

Default: 0 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Managing Subscriber Overcommitment

summaries

Syntax (summaries | no-summaries);

Hierarchy Level [edit logical-systems *logical-system-name* protocols (ospf | ospf3) **area** *area-id* **nssa**],
[edit logical-systems *logical-system-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **nssa**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf | ospf3) **area** *area-id* **nssa**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **nssa**],
[edit protocols (ospf | ospf3) **area** *area-id* **nssa**],
[edit protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)] area *area-id* **nssa**],
[edit routing-instances *routing-instance-name* protocols (ospf | ospf3) **area** *area-id* **nssa**],
[edit routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **nssa**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure whether or not area border routers advertise summary routes into an not-so-stubby area (NSSA):

- **summaries**—Flood summary link-state advertisements (LSAs) into the NSSA.
- **no-summaries**—Prevent area border routers from advertising summaries into an NSSA. If **default-metric** is configured for an NSSA, a Type 3 LSA is injected into the area by default.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **nssa on page 1642**
- **stub on page 1710**
- Configuring OSPF Areas

tag

Syntax tag *string*;

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options (**aggregate | generate | static**) (defaults | route)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options rib *routing-table-name* (**aggregate | generate | static**) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options (**aggregate | generate | static**) (defaults | route)],
[edit logical-systems *logical-system-name* routing-options rib *routing-table-name* (**aggregate | generate | static**) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options **aggregate | generate | static**) (defaults | route)],
[edit routing-instances *routing-instance-name* routing-options rib *routing-table-name* (**aggregate | generate | static**) (defaults | route)],
[edit routing-options (**aggregate | generate | static**) (defaults | route)],
[edit routing-options rib *routing-table-name* (**aggregate | generate | static**) (defaults | route)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Associate an OSPF tag with a static, aggregate, or generated route.

Options ***string***—OSPF tag string.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **aggregate on page 1470**
- **generate on page 1543**
- **static on page 1708**
- Configuring Static Routes
- Configuring Aggregate Routes
- Configuring Generated Routes

tcp-mss

Syntax tcp-mss *segment-size*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp group *group-name*],
[edit logical-systems *logical-system-name* protocols bgp group *group-name* neighbor *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *neighbor-name*],
[edit protocols bgp],
[edit protocol bgp group *group-name*],
[edit protocols bgp group *group-name* neighbor *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *neighbor-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the maximum segment size (MSS) for the TCP connection for BGP neighbors.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Limiting TCP Segment Size for BGP

threshold

Syntax threshold suppress *value* <reuse *value*>;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options **multicast** forwarding-cache],
[edit logical-systems *logical-system-name* routing-options **multicast** forwarding-cache],
[edit routing-instances *routing-instance-name* routing-options **multicast** forwarding-cache],
[edit routing-options **multicast** forwarding-cache]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the suppression and reuse thresholds for multicast forwarding cache limits.

Options

reuse *value*—Value at which to begin creating new multicast forwarding cache entries. This value is optional. If configured, this number should be less than the **suppress** value.

Range: 1 through 200,000

suppress *value*—Value at which to begin suppressing new multicast forwarding cache entries. This value is mandatory. This number should be greater than the **reuse** value.

Range: 1 through 200,000

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Multicast Forwarding Cache Limits

timeout (Flow Maps)

Syntax timeout (never | *minutes*);

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast **flow-map** *flow-map-name*],
[edit logical-systems *logical-system-name* routing-options multicast **flow-map** *flow-map-name*],
[edit routing-instances *routing-instance-name* routing-options multicast **flow-map** *flow-map-name*],
[edit routing-options multicast **flow-map** *flow-map-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the timeout value for multicast forwarding cache entries associated with the flow map.

Options ***minutes***—Length of time that the forwarding cache entry remains active.

Range: 1 through 720

never—Specify that the forwarding cache entry always remain active.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

timeout (Multicast)

Syntax timeout *minutes*;

Hierarchy Level [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast **forwarding-cache**],
[edit logical-systems *logical-system-name* routing-options multicast **forwarding-cache**],
[edit routing-instances *routing-instance-name* routing-options multicast **forwarding-cache**],
[edit routing-options multicast **forwarding-cache**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the timeout value for multicast forwarding cache entries.

Options ***minutes***—Length of time that the forwarding cache limit remains active.

Range: 1 through 720

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring General Multicast Forwarding Cache Properties

topologies

Syntax

```
topologies {
    ipv4-multicast;
    ipv6-multicast;
    ipv6-unicast;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols isis],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis],
[edit protocols isis],
[edit routing-instances *routing-instance-name* protocols isis]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure alternate IS-IS topologies.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring IS-IS Multicast Topologies

traceoptions (BGP)

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols bgp],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* protocols bgp **group** *group-name* neighbor *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*],
[edit protocols bgp],
[edit protocols bgp **group** *group-name*],
[edit protocols bgp group *group-name* **neighbor** *address*],
[edit routing-instances *routing-instance-name* protocols bgp],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* **neighbor** *address*]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure BGP protocol-level tracing options. To specify more than one tracing operation, include multiple flag statements.

Default

The default BGP protocol-level tracing options are inherited from the routing protocols **traceoptions** statement included at the **[edit routing-options]** hierarchy level. The default group-level trace options are inherited from the BGP protocol-level **traceoptions** statement. The default peer-level trace options are inherited from the group-level **traceoptions** statement.

Options

disable—(Optional) Disable the tracing operation. You can use this option is to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file *name*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place BGP tracing output in the file **bgp-log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 10 files

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

BGP Tracing Flags

- **4byte-as**—4-byte AS events
- **bfd**—BFD protocol events
- **damping**—Damping operations
- **graceful-restart**—Graceful restart events
- **keepalive**—BGP keepalive messages. If you enable the the BGP **update** flag only, received keepalive messages do not generate a trace message.
- **nsr-synchronization**—Nonstop routing synchronization events
- **open**—Open packets. These packets are sent between peers when they are establishing a connection.
- **packets**—All BGP protocol packets
- **refresh**—BGP refresh packets
- **update**—Update packets. These packets provide routing updates to BGP systems. If you enable only this flag, received keepalive messages do not generate a trace message. Use the **keepalive** flag to generate a trace message for keepalive messages.

Global Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information
- **filter**—Filter trace information. Applies only to **route** and **damping** tracing flags.
- **receive**—Packets being received.
- **send**—Packets being transmitted.

no-world-readable—(Optional) Prevent any user from reading the log file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*****.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*****.0** is renamed ***trace-file*****.1** and ***trace-file*** is renamed ***trace-file*****.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: ***x*****k** to specify KB, ***x*****m** to specify MB, or ***x*****g** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration.

Related Documentation

- **log-updown on page 1600** statement
- Tracing BGP Protocol Traffic
- Configuring OSPF Refresh and Flooding Reduction in Stable Topologies

traceoptions (IS-IS)

Syntax

```
traceoptions {
    file name <size size> <files number> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols **isis**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **isis**],
[edit protocols **isis**],
[edit routing-instances *routing-instance-name* protocols **isis**]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure IS-IS protocol-level tracing options. To specify more than one tracing operation, include multiple **flag** statements.

Default

The default IS-IS protocol-level tracing options are those inherited from the routing protocols **traceoptions** statement included at the **[edit routing-options]** hierarchy level.

Options

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file ***name***—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks (" "). All files are placed in the directory **/var/log**. We recommend that you place IS-IS tracing output in the file **isis-log**.

files ***number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 10 files

flag ***flag***—Tracing operation to perform. To specify more than one flag, include multiple **flag** statements.

IS-IS Protocol-Specific Tracing Flags

- **csn**—Complete sequence number PDU (CSNP) packets
- **error**—Errored IS-IS packets
- **graceful-restart**—Graceful restart operation
- **hello**—Hello packets

- **ldp-synchronization**—Synchronization between IS-IS and LDP
- **lsp**—Link-state PDU packets
- **lsp-generation**—Link-state PDU generation packets
- **packets**—All IS-IS protocol packets
- **psn**—Partial sequence number PDU (PSNP) packets
- **spf**—Shortest-path-first calculations

Global Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations, including adjacency changes

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

no-world-readable—(Optional) Prevent any user from reading the log file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

Note that if you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level
routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration.

Related Documentation
- Tracing IS-IS Protocol Traffic

traceoptions (OSPF)

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols (**ospf** | **ospf3**)],
[edit logical-systems *logical-system-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit protocols (**ospf** | **ospf3**)],
[edit protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)],
[edit routing-instances *routing-instance-name* protocols (**ospf** | **ospf3**)],
[edit routing-instances *routing-instance-name* protocols ospf3 **realm** (ipv4-unicast | ipv4-multicast | ipv6-multicast)]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure OSPF protocol-level tracing options.

To specify more than one tracing operation, include multiple **flag** statements.

Default

The default OSPF protocol-level tracing options are those inherited from the routing protocols **traceoptions** statement included at the **[edit routing-options]** hierarchy level.

Options

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place OSPF tracing output in the file **ospf-log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 10 files

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

OSPF Tracing Flags

- **database-description**—Database description packets, which are used in synchronizing the OSPF and OSPFv3 topological database.
- **error**—OSPF and OSPFv3 error packets.
- **event**—OSPF and OSPFv3 state transitions.
- **flooding**—Link-state flooding packets.
- **graceful-restart**—Graceful-restart events.
- **hello**—Hello packets, which are used to establish neighbor adjacencies and to determine whether neighbors are reachable.
- **ldp-synchronization**—Synchronization events between OSPF and LDP
- **lsa-ack**—Link-state acknowledgment packets, which are used in synchronizing the OSPF topological database.
- **lsa-analysis**—Link-state analysis packets
- **lsa-request**—Link-state request packets, which are used in synchronizing the OSPF topological database.
- **lsa-update**—Link-state updates packets, which are used in synchronizing the OSPF topological database.
- **nsr-synchronization**—Nonstop routing synchronization events.
- **on-demand**—Trace demand circuit extensions.
- **packet-dump**—Content of selected packet types.
- **packets**—All OSPF packets.
- **spf**—Shortest-path-first (SPF) calculations.

Global Tracing Flags

- **all**—All tracing operations.
- **general**—A combination of the **normal** and **route** trace operations.
- **normal**—All normal operations.

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions.
- **route**—Routing table changes.
- **state**—State transitions.
- **task**—Routing protocol task processing.
- **timer**—Routing protocol timer processing.

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information.
- **receive**—Packets being received.
- **send**—Packets being transmitted.

no-world-readable—(Optional) Prevent any user from reading the log file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration.

Related Documentation

- Tracing OSPF Protocol Traffic

traceoptions (RIP)

Syntax

```
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols **rip**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **rip**],
[edit protocols **rip**],
[edit routing-instances *routing-instance-name* protocols **rip**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set RIP protocol-level tracing options.

Default The default RIP protocol-level trace options are inherited from the global **traceoptions** statement.

Options

disable—(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file ***filename***—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. We recommend that you place RIP tracing output in the file **/var/log/rip-log**.

files ***number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum number of files, you must also specify a maximum file size with the size option.

Range: 2 through 1000 files

Default: 10 files

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

RIP Tracing Options

- **auth**—RIP authentication
- **error**—RIP error packets
- **expiration**—RIP route expiration processing
- **holddown**—RIP hold-down processing
- **nsr-synchronization**—Nonstop routing synchronization events
- **packets**—All RIP packets

- **request**—RIP information packets such as request, poll, and poll entry packets
- **trigger**—RIP triggered updates
- **update**—RIP update packets

Global Tracing Options

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information
- **receive**—Packets being received
- **receive-detail**—Provide detailed trace information for packets being received
- **send**—Packets being transmitted
- **send-detail**—Provide detailed trace information for packets being transmitted

no-world-readable—(Optional) Prevent any user from reading the log file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum file size, you must also specify a maximum number of trace files with the files option.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Tracing RIP Protocol Traffic

traceoptions (RIPng)

Syntax

```
traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable>;
    flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols ripng],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng],
[edit protocols ripng],
[edit routing-instances *routing-instance-name* protocols ripng]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set RIPng protocol-level tracing options.

Default The default RIPng protocol-level trace options are inherited from the global **traceoptions** statement.

Options

disable—(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. We recommend that you place RIPng tracing output in the file **/var/log/ripng-log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.
If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 10 files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

RIPng Tracing Options

- **error**—RIPng error packets
- **expiration**—RIPng route expiration processing
- **holddown**—RIPng hold-down processing
- **nsr-synchronization**—Nonstop routing synchronization events
- **packets**—All RIPng packets
- **request**—RIPng information packets such as request, poll, and poll entry packets

- **trigger**—RIPng triggered updates
- **update**—RIPng update packets

Global Tracing Options

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information
- **receive**—Packets being received
- **receive-detail**—Provide detailed trace information for packets being received
- **send**—Packets being transmitted
- **send-detail**—Provide detailed trace information for packets being transmitted

no-world-readable—(Optional) Do not allow any user to read the log file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Tracing RIPng Protocol Traffic

traceoptions (All Routing Protocols)

Syntax

```
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag <disable>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options],
[edit logical-systems *logical-system-name* routing-options],
[edit routing-instances *routing-instance-name* routing-options],
[edit routing-options],
[edit routing-options flow]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define tracing operations that track all routing protocol functionality in the routing device.

To specify more than one tracing operation, include multiple **flag** statements.

Default If you do not include this statement, no global tracing operations are performed.

Options **Values:**

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place global routing protocol tracing output in the file **routing-log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 10 files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. These are the global routing protocol tracing options:

- **all**—All tracing operations
- **condition-manager**—Condition-manager events
- **config-internal**—Configuration internals

- **general**—All normal operations and routing table changes (a combination of the **normal** and **route** trace operations)
- **graceful-restart**—Graceful restart operations
- **normal**—All normal operations
- **nsr-synchronization**—Nonstop active routing synchronization
- **parse**—Configuration parsing
- **policy**—Routing policy operations and actions
- **regex-parse**—Regular-expression parsing
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

no-world-readable—(Optional) Prevent any user from reading the log file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. Note that if you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration.

Related Documentation

- Tracing Global Routing Protocol Operations

traffic-engineering (OSPF)

Syntax

```
traffic-engineering {
  <advertise-unnumbered-interfaces>;
  <credibility-protocol-preference>;
  ignore-lsp-metrics;
  multicast-rpf-routes;
  no-topology;
  shortcuts {
    lsp-metric-into-summary;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols (**ospf** | ospf3)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (**ospf** | ospf3)],
[edit protocols (**ospf** | ospf3)],
[edit routing-instances *routing-instance-name* protocols (**ospf** | ospf3)]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Enable the OSPF traffic engineering features.

Default

Traffic engineering support is disabled.

Options

advertise-unnumbered-interfaces—(Optional) (OSPFv2 only) Include the link-local identifier in the link-local traffic-engineering link-state advertisement. You do not need to include this statement if RSVP is able to signal unnumbered interfaces as defined in RFC 3477.

credibility-protocol-preference—(Optional) (OSPFv2 only) Specify to use the configured preference value for OSPF routes to calculate the traffic engineering database credibility value used to select IGP routes. Use this statement to override the default behavior of having the traffic engineering database prefer IS-IS routes even if OSPF routes are configured a with a lower, that is, preferred, preference value.

multicast-rpf-routes—(Optional) (OSPFv2 only) Install routes for multicast RPF checks into the **inet.2** routing table.

no-topology—(Optional) (OSPFv2 only) Disable the dissemination of the link-state topology information.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Enabling OSPF Traffic Engineering Support

transit-delay

Syntax transit-delay *seconds*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols ospf area *area-id* peer-interface *interface-name*],
[edit logical-systems *logical-system-name* protocols (ospf | ospf3) area *area-id* **interface** *interface-name*],
[edit logical-systems *logical-system-name* protocols (ospf | ospf3) area *area-id* **virtual-link**],
[edit logical-systems *logical-system-name* protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf area *area-id* **interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf area *area-id* **virtual-link**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **interface** *interface-name*],
[edit protocols ospf area *area-id* peer-interface *interface-name*],
[edit protocols (ospf | ospf3) area *area-id* **interface** *interface-name*],
[edit protocols (ospf | ospf3) area *area-id* **virtual-link**],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast)] area *area-id* **interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols ospf area *area-id* **interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols ospf area *area-id* **virtual-link**],
[edit routing-instances *routing-instance-name* protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* **interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set the estimated time required to transmit a link-state update on the interface. When calculating this time, make sure to account for transmission and propagation delays.

You should never have to modify the transit delay time.

Options ***seconds***—Estimated time, in seconds.

Range: 1 through 65,535 seconds

Default: 1 second

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring OSPF Timers

type

Syntax type *type*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols bgp **group** *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols bgp **group** *group-name*],
[edit protocols bgp **group** *group-name*],
[edit routing-instances *routing-instance-name* protocols bgp **group** *group-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the type of BGP peer group.

Options ***type***—Type of group:

- **external**—External group
- **internal**—Internal group

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring BGP Groups and Peers

type-7

Syntax type-7;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols (ospf | ospf3) area *area-id* nssa **default-lsa**],
[edit logical-systems *logical-system-name* protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* nssa **default-lsa**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* nssa **default-lsa**],
[edit logical-systems *logical-system-name* protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* nssa **default-lsa**],
[edit protocols (ospf | ospf3) area *area-id* nssa **default-lsa**],
[edit protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* nssa **default-lsa**],
[edit routing-instances *routing-instance-name* protocols (ospf | ospf3) area *area-id* nssa **default-lsa**],
[edit routing-instances *routing-instance-name* protocols ospf3 realm (ipv4-unicast | ipv4-multicast | ipv6-multicast) area *area-id* nssa **default-lsa**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Flood Type 7 default link-state advertisements (LSAs) if the **no-summaries** statement is configured.

By default, when the **no-summaries** statement is configured, a Type 3 LSA is injected into not-so-stubby areas (NSSAs). This statement enables NSSA ABRs to advertise a Type 7 default LSA into the NSSA if you have also included the **no-summaries** statement in the configuration.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring OSPF Areas

update-interval

Syntax update-interval *seconds*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols **rip**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **rip**],
[edit protocols **rip**],
[edit routing-instances *routing-instance-name* protocols **rip**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure an update time interval to periodically send out routes learned by RIP to neighbors.

Options ***seconds***—Estimated time to wait before making updates to the routing table.

Range: 10 through 60 seconds

Default: 30 seconds

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring RIP Timers

update-interval

Syntax update-interval *seconds*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols ripng],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng],
[edit protocols ripng],
[edit routing-instances *routing-instance-name* protocols ripng]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure an update time interval to periodically send out routes learned by RIP to neighbors.

Options ***seconds***—Estimated time to wait before making updates to the routing table.

Range: 10 through 60 seconds

Default: 30 seconds

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring RIP Timers

upstream-interface

Syntax upstream-interface [*interface-names*];

Hierarchy Level
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast **pim-to-igmp-proxy**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* routing-options multicast **pim-to-mld-proxy**],
[edit logical-systems *logical-system-name* routing-options multicast **pim-to-igmp-proxy**],
[edit logical-systems *logical-system-name* routing-options multicast **pim-to-mld-proxy**],
[edit routing-instances *routing-instance-name* routing-options multicast **pim-to-igmp-proxy**],
[edit routing-instances *routing-instance-name* routing-options multicast **pim-to-mld-proxy**],
[edit routing-options multicast **pim-to-igmp-proxy**],
[edit routing-options multicast **pim-to-mld-proxy**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure at least one, but not more than two, upstream interfaces on the rendezvous point (RP) routing device that resides between a customer edge–facing Protocol Independent Multicast (PIM) domain and a core-facing PIM domain. The RP routing device translates PIM join or prune messages into corresponding IGMP report or leave messages (if you include the **pim-to-igmp-proxy** statement), or into corresponding MLD report or leave messages (if you include the **pim-to-mld-proxy** statement). The routing device then proxies the IGMP or MLD report or leave messages to one or both upstream interfaces to forward IPv4 multicast traffic (for IGMP) or IPv6 multicast traffic (for MLD) across the PIM domains.

Options
interface-names—Names of one or two upstream interfaces to which the RP routing device proxies IGMP or MLD report or leave messages for transmission of multicast traffic across PIM domains. You can specify a maximum of two upstream interfaces on the RP routing device. To configure a set of two upstream interfaces, specify the full interface names, including all physical and logical address components, within square brackets ([ ]). For details about specifying interfaces, see the *Junos OS Network Interfaces Configuration Guide*.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring PIM-to-IGMP Message Translation
- Configuring PIM-to-MLD Message Translation

virtual-inet6-address

Syntax virtual-inet6-address [*addresses*];

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* vrrp-inet6-group *group-id*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the addresses of the virtual routers in a Virtual Router Redundancy Protocol (VRRP) IPv6 group. You can configure up to eight addresses.

NOTE: The address of an aggregated Ethernet interface (a LAG) or a routed VLAN interface (RVI) cannot be assigned as the virtual router address in a VRRP IPv6 group.

Options ***addresses***—Addresses of one or more virtual routers. Do not include a prefix length. If the address is the same as the interface's physical address, the interface becomes the master virtual router for the group.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring VRRP for IPv6 (CLI Procedure) on page 1452

virtual-link

Syntax

```
virtual-link neighbor-id router-id transit-area area-id {
    disable;
    authentication key <key-id identifier>;
    dead-interval seconds;
    hello-interval seconds;
    ipsec-sa name;
    retransmit-interval seconds;
    transit-delay seconds;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols (ospf | ospf3) area *area-id*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ospf area *area-id*],
[edit protocols (ospf | ospf3) area *area-id*],
[edit routing-instances *routing-instance-name* protocols ospf area *area-id*]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

For backbone areas only, create a virtual link to use in place of an actual physical link. All area border routers and other routing devices on the backbone must be contiguous. If this is not possible and there is a break in OSPF connectivity, use virtual links to create connectivity to the OSPF backbone. When configuring virtual links, you must configure links on the two routing devices that form the end points of the link, and both these two routing devices must be area border routers. You cannot configure links through stub areas.

Options

neighbor-id *router-id*—IP address of the routing device at the remote end of the virtual link.

transit-area *area-id*—Area identifier of the area through which the virtual link transits. Virtual links are not allowed to transit the backbone area.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring OSPF Areas

virtual-link-local-address

Syntax virtual-link-local-address *ipv6-address*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet6 address *address* vrrp-inet6-group *group-id*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a virtual link local address for a Virtual Router Redundancy Protocol (VRRP) IPv6 group. You must explicitly define a virtual link local address for each VRRP IPv6 group. The virtual link local address must be in the same subnet as the physical interface address.

Options ***ipv6-address***—Virtual link local IPv6 address for VRRP for an IPv6 group.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring VRRP for IPv6 (CLI Procedure) on page 1452

vrrp-inet6-group

Syntax

```
vrrp-inet6-group group-id {
  inet6-advertise-interval milliseconds;
    preempt{
      hold-time seconds;
    }
    priority number;
    virtual-inet6-address;
    virtual-link-local-address
  }
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet6 address *address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a Virtual Router Redundancy Protocol (VRRP) IPv6 group.

Options ***group-id***—VRRP group identifier. If you enable MAC source address filtering on the interface, you must include the virtual MAC address in the list of source MAC addresses that you specify in the **source-address-filter** statement. MAC addresses ranging from **00:00:5e:00:01:00** through **00:00:5e:00:01:ff** are reserved for VRRP, as defined in RFC 3768. The VRRP group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.

Range: 0 through 255

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring VRRP for IPv6 (CLI Procedure) on page 1452

wide-metrics-only

Syntax wide-metrics-only;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols isis **level** *level-number*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols isis **level** *level-number*],
[edit protocols isis **level** *level-number*],
[edit routing-instances *routing-instance-name* protocols isis **level** *level-number*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure IS-IS to generate metric values greater than 63 on a per IS-IS level basis.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- te-metric
- Enabling Wide IS-IS Metrics for Traffic Engineering

CHAPTER 74

Operational Commands for Layer 3 Protocols

clear (ospf | ospf3) database

Syntax

clear (ospf | ospf3) database
<advertising-router (*router-id* | self)>
<area *area-id*>
<asbrsummary>
<external>
<instance *instance-name*>
<inter-area-prefix>
<inter-area-router>
<intra-area-prefix>
<link-local>
<logical-system (all | *logical-system-name*)>
<lsa-id *lsa-id*>
<netsummary>
<network>
<nssa>
<opaque-area>
<purge>
<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>
<router>

Syntax (J-EX Series Switch)

clear (ospf | ospf3) database
<advertising-router (*router-id* | self)>
<area *area-id*>
<asbrsummary>
<external>
<instance *instance-name*>
<inter-area-prefix>
<inter-area-router>
<intra-area-prefix>
<link-local>
<lsa-id *lsa-id*>
<netsummary>
<network>
<nssa>
<opaque-area>
<purge>
<router>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

With the master Routing Engine, delete entries in the Open Shortest Path First (OSPF) link-state advertisement (LSA) database. With the backup Routing Engine, delete the OSPF LSA database and sync the new database with the master Routing Engine. You can also use the **purge** command with any of the options to discard rather than delete the specified LSA entries.

CAUTION: This command is useful only for testing. Use it with care, because it causes significant network disruption.

Options

none—Delete all LSAs other than the system's own LSAs, which are regenerated. To resynchronize the database, the system destroys all adjacent neighbors that are in the state **EXSTART** or higher. The neighbors are then reacquired and the databases are synchronized.

advertising-router (*router-id* | self)—(Optional) Discard entries for the LSA entries advertised by the specified routing device or by this routing device.

area *area-id*—(Optional) Discard entries for the LSAs in the specified area.

asbrsummary—(Optional) Discard summary AS boundary router LSA entries.

external—(Optional) Discard external LSAs.

instance *instance-name*—(Optional) Delete or discard entries for the specified routing instance only.

inter-area-prefix—(OSPFv3 only) (Optional) Discard interarea prefix LSAs.

inter-area-router—(OSPFv3 only) (Optional) Discard interarea router LSAs.

intra-area-prefix—(OSPFv3 only) (Optional) Discard intra-area prefix LSAs.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

link-local—(Optional) Delete link-local LSAs.

lsa-id *lsa-id*—(Optional) Discard the LSA entries with the specified LSA identifier.

netsummary—(Optional) Discard summary network LSAs.

network—(Optional) Discard network LSAs.

nssa—(Optional) Discard not-so-stubby area (NSSA) LSAs.

opaque-area—(Optional) Discard opaque area-scope LSAs.

realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)—(OSPFv3 only) (Optional) Delete the entries for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

router—(Optional) Discard router LSAs.

purge—(Optional) Discard all entries in the link-state advertisement database. All link-state advertisements are set to **MAXAGE** and are flooded. The database is repopulated when the originators of the link-state advertisements receive the **MAXAGE** link-state advertisements and reissue them.

Required Privilege Level

clear

Related Documentation

- **show ospf database on page 1878**

- show ospf3 database on page 1868

List of Sample Output clear ospf database on page 1750

Output Fields When you enter this command, you are provided feedback on the status of your request.

clear ospf database

```
user@host> clear ospf database
```

clear (ospf | ospf3) io-statistics

Syntax clear (ospf | osfp3) statistics
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch) clear (ospf | osfp3) statistics

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Clear Open Shortest Path First (OSPF) input and output statistics.

Options none—Clear OSPF input and output statistics.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level clear

List of Sample Output **clear ospf io-statistics on page 1751**

Output Fields When you enter this command, you are provided feedback on the status of your request.

clear ospf io-statistics

```
user@host> clear ospf io-statistics
```

clear (ospf | ospf3) neighbor

Syntax

clear (ospf | ospf3) neighbor
<area *area-id*>
<instance *instance-name*>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>
<neighbor>
<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>

Syntax (J-EX Series Switch)

clear (ospf | ospf3) neighbor
<area *area-id*>
<instance *instance-name*>
<interface *interface-name*>
<neighbor>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Tear down Open Shortest Path First (OSPF) neighbor connections.

Options

none—Tear down OSPF connections with all neighbors for all routing instances.

area *area-id*—(Optional) Tear down neighbor connections for the specified area only.

instance *instance-name*—(Optional) Tear down neighbor connections for the specified routing instance only.

interface *interface-name*—(Optional) Tear down neighbor connections for the specified interface only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

neighbor—(Optional) Clear the state of the specified neighbor only.

realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)—(Optional) (OSPFv3 only) Clear the state of the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

Required Privilege Level

clear

Related Documentation

- **show (ospf | ospf3) neighbor on page 1782**

List of Sample Output

clear ospf neighbor on page 1752

Output Fields

When you enter this command, you are provided feedback on the status of your request.

clear ospf neighbor

```
user@host> clear ospf neighbor
```

clear (ospf | ospf3) statistics

Syntax
clear (ospf | osfp3) statistics
<logical-system (all | *logical-system-name*)>
<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>

Syntax (J-EX Series Switch)
clear (ospf | osfp3) statistics

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Clear Open Shortest Path First (OSPF) statistics.

Options
none—Clear OSPF statistics.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)—(Optional) (OSPFv3 only) Clear statistics for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

Required Privilege Level
clear

Related Documentation
- **show (ospf | ospf3) statistics on page 1796**

List of Sample Output
clear ospf statistics on page 1753

Output Fields
See **show (ospf | ospf3) statistics** for an explanation of output fields.

clear ospf statistics
The following sample output displays OSPF statistics before and after the **clear ospf statistics** command is entered:

```
user@host> show ospf statistics

Packet type             Total                 Last 5 seconds
                    Sent      Received        Sent      Received
   Hello            3254          2268           3             1
     DbD              41            46           0             0
   LSReq               8             7           0             0
LSUpdate             212           154           0             0
   LSAck              65            98           0             0

LSAs retransmitted: 3, last 5 seconds: 0

Flood queue depth: 0
Total rexmit entries: 0, db summaries: 0, lsreq entries: 0

Receive errors:
  626 subnet mismatches


user@host> clear ospf statistics
```

```
user@host> show ospf statistics
Packet type             Total                 Last 5 seconds
                    Sent      Received        Sent      Received
   Hello               3             1           3             1
     DbD               0             0           0             0
   LSReq               0             0           0             0
LSUpdate               0             0           0             0
   LSAck               0             0           0             0

LSAs retransmitted: 0, last 5 seconds: 0

Flood queue depth: 0
Total rexmit entries: 0, db summaries: 0, lsreq entries: 0
Receive errors:
  None
```

clear bgp damping

Syntax

clear bgp damping
<logical-system (all | *logical-system-name*)>
<*prefix*>

Syntax (J-EX Series Switch)

clear bgp damping
<*prefix*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Clear Border Gateway Protocol (BGP) route flap damping information.

Options

none—Clear all BGP route flap damping information.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

prefix—(Optional) Clear route flap damping information for only the specified destination prefix.

Required Privilege Level

clear

Related Documentation

- **show policy damping on page 1886**
- **show route damping on page 1922**

List of Sample Output

clear bgp damping on page 1755

Output Fields

When you enter this command, you are provided feedback on the status of your request.

clear bgp damping

```
user@host> clear bgp damping
```

clear bgp neighbor

Syntax

clear bgp neighbor
<as *as-number*>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<*neighbor*>
<soft | soft-inbound>
<soft-minimum-igp>

Syntax (J-EX Series Switch)

clear bgp neighbor
<as *as-number*>
<instance *instance-name*>
<*neighbor*>
<soft | soft-inbound>
<soft-minimum-igp>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Perform one of the following tasks:

- Change the state of one or more Border Gateway Protocol (BGP) neighbors to **IDLE**. For neighbors in the **ESTABLISHED** state, this command drops the TCP connection to the neighbors and then reestablishes the connection.
- (**soft** or **soft-inbound** keyword only) Reapply export policies or import policies, respectively, and send refresh updates to one or more BGP neighbors without changing their state.

Options

none—Change the state of all BGP neighbors to **IDLE**.

as *as-number*—(Optional) Apply this command only to neighbors in the specified autonomous system (AS).

instance *instance-name*—(Optional) Apply this command only to neighbors for the specified routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

neighbor—(Optional) IP address of a BGP peer. Apply this command only to the specified neighbor.

soft—(Optional) Reapply any export policies and send refresh updates to neighbors without clearing the state.

soft-inbound—(Optional) Reapply any import policies and send refresh updates to neighbors without clearing the state.

soft-minimum-igp—(Optional) Provides soft refresh of the outbound state when the interior gateway protocol (IGP) metric is reset.

Required Privilege Level clear

Related Documentation
- show bgp neighbor on page 1812

List of Sample Output clear bgp neighbor on page 1757

Output Fields When you enter this command, you are provided feedback on the status of your request.

clear bgp neighbor

```
user@host> clear bgp neighbor
```

clear bgp table

Syntax **clear bgp table *table-name***
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch) **clear bgp table *table-name***

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Request BGP to refresh routes in a specified routing table.

Options logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

table-name—Request that BGP refresh routes in the specified table.

Additional Information In some cases, a prefix limit is associated with a routing table for a VPN instance. When this limit is exceeded (for example, because of a network misconfiguration), some routes might not be inserted in the table. Such routes need to be added to the table after the network issue is resolved. Use the **clear bgp table** command to request that BGP refresh routes in a VPN instance table.

Required Privilege Level clear

List of Sample Output **clear bgp table private.inet.0 on page 1758**

Output Fields This command produces no output.

clear bgp table private.inet.0

```
user@host> clear bgp table private.inet.0
```

clear ipv6 neighbors

Syntax clear ipv6 neighbors
<all | host *hostname*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Clear IPv6 neighbor cache information.

Options none—Clear all IPv6 neighbor cache information.

all—(Optional) Clear all IPv6 neighbor cache information.

host *hostname*—(Optional) Clear the information for the specified IPv6 neighbors.

Required Privilege Level view

Related Documentation
- **show ipv6 neighbors on page 1031**

List of Sample Output **clear ipv6 neighbors on page 1759**

Output Fields When you enter this command, you are provided feedback on the status of your request.

clear ipv6 neighbors

```
user@host> clear ipv6 neighbors
```

clear isis adjacency

Syntax

clear isis adjacency
<instance *instance-name*>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>
<*neighbor*>

Syntax (J-EX Series Switch)

clear isis adjacency
<instance *instance-name*>
<interface *interface-name*>
<*neighbor*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Remove entries from the Intermediate System-to-Intermediate System (IS-IS) adjacency database.

Options

none—Remove all entries from the adjacency database.

instance *instance-name*—(Optional) Clear all adjacencies for the specified routing instance only.

interface *interface-name*—(Optional) Clear all adjacencies for the specified interface only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

neighbor—(Optional) Clear adjacencies for the specified neighbor only.

Required Privilege Level

clear

Related Documentation

- **show isis adjacency on page 1830**

List of Sample Output

clear isis adjacency on page 1760

Output Fields

See **show isis adjacency** for an explanation of output fields.

clear isis adjacency

The following sample output displays IS-IS adjacency database information before and after the **clear isis adjacency** command is entered:

```
user@host> show isis adjacency
IS-IS adjacency database:
Interface     System         L State       Hold (secs) SNPA
so-1/0/0.0    karakul        3 Up                   26
so-1/1/3.0    1921.6800.5080 3 Up                   23
so-5/0/0.0    1921.6800.5080 3 Up                   19

user@host> clear isis adjacency karakul
```

```
user@host> show isis adjacency
IS-IS adjacency database:
Interface     System         L State        Hold (secs) SNPA
so-1/0/0.0    karakul        3 Initializing          26
so-1/1/3.0    1921.6800.5080 3 Up                    24
so-5/0/0.0    1921.6800.5080 3 Up                    21
```

clear isis database

Syntax

clear isis database
<*entries*>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<purge>

Syntax (J-EX Series Switch)

clear isis database
<*entries*>
<instance *instance-name*>
<purge>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Remove the entries from the Intermediate System-to-Intermediate System (IS-IS) link-state database, which contains prefixes and topology information. You can also use **purge** with any of the options to initiate a network-wide purge of link-state PDUs (LSPs) rather than the local deletion of entries from the IS-IS link-state database.

CAUTION: In a production network, the purge command option may cause short-term network-wide traffic disruptions. Use with caution!

Options

none—Remove all entries from the IS-IS link-state database for all routing instances.

entries—(Optional) Name of the database entry.

instance *instance-name*—(Optional) Clear all entries for the specified routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

purge—(Optional) Discard all entries in the IS-IS link-state database.

Required Privilege Level

clear

Related Documentation

- **show isis database on page 1843**

List of Sample Output

clear isis database on page 1762

Output Fields

See **show isis database** for an explanation of output fields.

clear isis database

The following sample output displays IS-IS link-state database information before and after the **clear isis database** command is entered:

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                 Sequence Checksum Lifetime (secs)
crater.00-00               0x12   0x84dd            1139
```

```
  1 LSPs
IS-IS level 2 link-state database:
LSP ID                 Sequence Checksum Lifetime (secs)
crater.00-00               0x19   0xe92c            1134
badlands.00-00             0x16   0x1454             985
carlsbad.00-00             0x33   0x220b            1015
ranier.00-00               0x2e   0xfc31            1007
1921.6800.5066.00-00       0x11   0x7313             566
1921.6800.5067.00-00       0x14   0xd9d4             939
  6 LSPs

user@host> clear isis database

user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                 Sequence Checksum Lifetime (secs)

IS-IS level 2 link-state database:
LSP ID                 Sequence Checksum Lifetime (secs)
```

clear isis overload

Syntax
clear isis overload
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
clear isis overload
<instance *instance-name*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Reset the Intermediate System-to-Intermediate System (IS-IS) dynamic overload bit. This command can appear to not work, continuing to display **overload** after execution. The bit is reset only if the root cause is corrected by configuration remotely or locally.

Options
none—Reset the IS-IS dynamic overload bit.

instance *instance-name*—(Optional) Reset the IS-IS dynamic overload bit for the specified routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
clear

Related Documentation
- **show isis database on page 1843**

List of Sample Output
clear isis overload on page 1764

Output Fields
See **show isis database** for an explanation of output fields.

clear isis overload
The following sample output displays IS-IS database information before and after the **clear isis overload** command is entered:

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                  Sequence Checksum Lifetime Attributes
pro3-c.00-00                 0x4   0x10db     1185 L1 L2 Overload

  1 LSPs
IS-IS level 2 link-state database:
LSP ID                  Sequence Checksum Lifetime Attributes
pro3-c.00-00                 0x5   0x429f     1185 L1 L2 Overload


pro2-a.00-00                0x91e   0x2589      874 L1 L2
pro2-a.02-00                  0x1    0xcbc      874 L1 L2
  3 LSPs

user@host> clear isis overload
```

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                  Sequence Checksum Lifetime Attributes
pro3-c.00-00                0xa   0x429e     1183 L1 L2
  1 LSPs

IS-IS level 2 link-state database:
LSP ID                  Sequence Checksum Lifetime Attributes
pro3-c.00-00                0xc   0x9c39     1183 L1 L2
pro2-a.00-00              0x91e   0x2589      783 L1 L2
pro2-a.02-00                0x1    0xcbc      783 L1 L2
  3 LSPs
```

clear isis statistics

Syntax
clear isis statistics
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
clear isis statistics
<instance *instance-name*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Set statistics about Intermediate System-to-Intermediate System (IS-IS) traffic to zero.

Options
none—Set IS-IS traffic statistics to zero for all routing instances.

instance *instance-name*—(Optional) Set IS-IS traffic statistics to zero for the specified routing instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- **show isis statistics on page 1866**

List of Sample Output
clear isis statistics on page 1766

Output Fields
See **show isis statistics** for an explanation of output fields.

clear isis statistics
The following sample output displays IS-IS statistics before and after the **clear isis statistics** command is entered:

```
user@host> show isis statistics
IS-IS statistics for merino:

PDU type     Received  Processed      Drops       Sent     Rexmit
LSP             12793      12793          0       8666        719
IIH            116751     116751          0     118834          0
CSNP           203956     203956          0     204080          0
PSNP             7356       7350          6       8635          0
Unknown             0          0          0          0          0
Totals         340856     340850          6     340215        719

Total packets received: 340856 Sent: 340934

SNP queue length:           0 Drops:          0
LSP queue length:           0 Drops:          0

SPF runs:                1064
Fragments rebuilt:       1087
LSP regenerations:        436
```

```
Purges initiated:           0

user@host> clear isis statistics

user@host> show isis statistics
IS-IS statistics for merino:

PDU type     Received  Processed      Drops       Sent     Rexmit
LSP                 0          0          0          0          0
IIH                 3          3          0          3          0
CSNP                2          2          0          4          0
PSNP                0          0          0          0          0
Unknown             0          0          0          0          0
Totals              5          5          0          7          0

Total packets received: 5 Sent: 7

SNP queue length:           0 Drops:          0
LSP queue length:           0 Drops:          0

SPF runs:                   0
Fragments rebuilt:          0
LSP regenerations:          0
Purges initiated:           0
```

clear ospf overload

Syntax
clear ospf overload
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
clear ospf overload
<instance *instance-name*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Clear the Open Shortest Path First version 2 (OSPFv2) overload bit and rebuild link-state advertisements (LSAs).

Options
none—Clear the overload bit and rebuild LSAs for all routing instances.

instance *instance-name*—(Optional) Clear the overload bit and rebuild LSAs for the specified routing instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
clear

List of Sample Output
clear ospf overload on page 1768

Output Fields
When you enter this command, you are provided feedback on the status of your request.

clear ospf overload
```
user@host> clear ospf overload
```

clear rip general-statistics

Syntax
clear rip general-statistics
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
clear rip general-statistics

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Clear Routing Information Protocol (RIP) general statistics.

Options
none—Clear RIP general statistics.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
clear

Related Documentation
- **show rip general-statistics on page 1888**

List of Sample Output
clear rip general-statistics on page 1769

Output Fields
When you enter this command, you are provided feedback on the status of your request.

clear rip general-statistics
```
user@host> clear rip general-statistics
```

clear rip statistics

Syntax

clear rip statistics
<instance (all | *instance-name*)>
<logical-system (all | *logical-system-name*)>
<*neighbor*>

Syntax (J-EX Series Switch)

clear rip statistics
<instance (all | *instance-name*)>
<*neighbor*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Clear Routing Information Protocol (RIP) statistics.

Options

none—Reset RIP counters for all neighbors for all routing instances.

instance (all | *instance-name*)—(Optional) Clear RIP statistics for all instances or for the specified routing instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

neighbor—(Optional) Clear RIP statistics for the specified neighbor only.

Required Privilege Level

clear

Related Documentation

- **show rip statistics on page 1891**

List of Sample Output

clear rip statistics on page 1770

Output Fields

When you enter this command, you are provided feedback on the status of your request.

clear rip statistics

```
user@host> clear rip statistics
```

clear ripng general-statistics

Syntax
clear ripng general-statistics
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
clear ripng general-statistics

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Clear Routing Information Protocol next generation (RIPng) general statistics.

Options
none—Clear RIPng general statistics.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
clear

Related Documentation
- **show ripng general-statistics on page 1894**

List of Sample Output
clear ripng general-statistics on page 1771

Output Fields
When you enter this command, you are provided feedback on the status of your request.

clear ripng general-statistics
```
user@host> clear ripng general-statistics
```

clear ripng statistics

Syntax
clear ripng statistics
<*instance* | *name*>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
clear ripng statistics
<*instance* | *name*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Clear Routing Information Protocol next-generation (RIPng) statistics.

Options
none—Reset RIPng counters for all neighbors for all routing instances.

instance—(Optional) Reset RIPng counters for the specified instance.

name—(Optional) Reset RIPng counters for the specified neighbor.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
clear

Related Documentation
- **show ripng statistics on page 1897**

List of Sample Output
clear ripng statistics on page 1772

Output Fields
When you enter this command, you are provided feedback on the status of your request.

clear ripng statistics
```
user@host> clear ripng statistics
```

show (ospf | ospf3) interface

Syntax

show (ospf | ospf3) interface
<brief | detail | extensive>
<area *area-id*>
<*interface-name*>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<realm (ip4-multicast | ipv4-unicast | ipv6-multicast)>

Syntax (J-EX Series Switch)

show (ospf | ospf3) interface
<brief | detail | extensive>
<area *area-id*>
<*interface-name*>
<instance *instance-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display the status of Open Shortest Path First (OSPF) interfaces.

Options

none—Display standard information about the status of all OSPF interfaces for all routing instances

brief | detail | extensive—(Optional) Display the specified level of output.

area *area-id*—(Optional) Display information about the interfaces that belong to the specified area.

interface-name—(Optional) Display information for the specified interface.

instance *instance-name*—(Optional) Display all OSPF interfaces under the named routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)—(Optional) (OSPFv3 only) Display information about the interfaces for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

Required Privilege Level

view

List of Sample Output

show ospf interface brief on page 1775
show ospf interface detail on page 1775
show ospf3 interface detail on page 1776
show ospf interface detail(When Multiarea Adjacency Is Configured) on page 1776
show ospf interface area area-id on page 1777
show ospf interface extensive (When Flooding Reduction Is Enabled) on page 1777

Output Fields Table 203 on page 1774 lists the output fields for the **show (ospf | ospf3) interface** command. Output fields are listed in the approximate order in which they appear.

Table 203: show (ospf | ospf3) interface Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface running OSPF version 2 or OSPF version 3.	All levels
State	State of the interface: **BDR**, **Down**, **DR**, **DRother**, **Loop**, **PtToPt**, or **Waiting**.	All levels
Area	Number of the area that the interface is in.	All levels
DR ID	Address of the area's designated router.	All levels
BDR ID	Backup designated router for a particular subnet.	All levels
Nbrs	Number of neighbors on this interface.	All levels
Type	Type of interface: **LAN**, **NBMA**, **P2MP**, **P2P**, or **Virtual**.	**detail extensive**
Address	IP address of the neighbor.	**detail extensive**
Mask	Netmask of the neighbor.	**detail extensive**
Prefix-length	(OSPFv3) IPv6 prefix length, in bits.	**detail extensive**
OSPF3-Intf-Index	(OSPFv3) OSPF version 3 interface index.	**detail extensive**
MTU	Interface's maximum transmission unit (MTU).	**detail extensive**
Cost	Interface's cost (metric).	**detail extensive**
DR addr	Address of the designated router.	**detail extensive**
BDR addr	Address of the backup designated router.	**detail extensive**
Adj count	Number of adjacent neighbors.	**detail extensive**
Secondary	Indicates that this interface is configured as a secondary interface for this area. This interface can belong to more than one area, but can be designated as a primary interface only for one area.	**detail extensive**
Flood Reduction	Indicates that this interface is configured with flooding reduction. All self-originated LSAs from this interface are initially sent with the **DoNotAge** bit set. As a result, LSAs are refreshed only when a change occurs.	**extensive**
Priority	Router priority used in designated router (DR) election on this interface.	**detail extensive**
Flood list	List of link-state advertisements (LSAs) that might be about to flood this interface.	**extensive**

Table 203: show (ospf | ospf3) interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
Ack list	Acknowledgment list. List of pending acknowledgments on this interface.	**extensive**
Descriptor list	List of packet descriptors.	**extensive**
Hello	Configured value for the Hello timer.	**detail extensive**
Dead	Configured value for the Dead timer.	**detail extensive**
Auth type	(OSPFv2) Authentication mechanism for sending and receiving OSPF protocol packets: • **MD5**—MD5 mechanism is configured in accordance with RFC 2328. • **None**—No authentication method is configured. • **Password**—Simple password (RFC 2328) is configured.	**detail extensive**
Topology	(Multiarea adjacency) Name of topology: **default** or ***name***	
IPSec SA name	(OSPFv2) Name of the IPsec security association name	**detail extensive**
Active key ID	(OSPFv2 and MD5) Number from **0** to **255** that uniquely identifies an MD5 key.	**detail extensive**
Start time	(OSPFv2 and MD5) Time at which the routing device starts using an MD5 key to authenticate OSPF packets transmitted on the interface on which this key is configured. To authenticate received OSPF protocol packets, the key becomes effective immediately after the configuration is committed. If the start time option is not configured, the key is effective immediately for send and receive and is displayed as **Start time 1970 Jan 01 00:00:00 PST.**	**detail extensive**
ReXmit	Configured value for the Retransmit timer.	**detail extensive**
Stub, **Not Stub**, or **Stub NSSA**	Type of area.	**detail extensive**

show ospf interface brief

```
user@host> show ospf interface brief
Intf                State     Area            DR ID           BDR ID       Nbrs
at-5/1/0.0          PtToPt    0.0.0.0         0.0.0.0         0.0.0.0         1
ge-2/3/0.0          DR        0.0.0.0         192.168.4.16    192.168.4.15    1
lo0.0               DR        0.0.0.0         192.168.4.16    0.0.0.0         0
so-0/0/0.0          Down      0.0.0.0         0.0.0.0         0.0.0.0         0
so-6/0/1.0          PtToPt    0.0.0.0         0.0.0.0         0.0.0.0         1
so-6/0/2.0          Down      0.0.0.0         0.0.0.0         0.0.0.0         0
so-6/0/3.0          PtToPt    0.0.0.0         0.0.0.0         0.0.0.0         1
```

show ospf interface detail

```
user@host> show ospf interface detail
Interface             State     Area            DR ID           BDR ID Nbrs
fe-0/0/1.0            BDR      0.0.0.0         192.168.37.12   10.255.245.215 1
Type LAN, address 192.168.37.11, Mask 255.255.255.248, MTU 4460, Cost 40
DR addr 192.168.37.12, BDR addr 192.168.37.11, Adj count 1, Priority 128
Hello 10, Dead 40, ReXmit 5, Not Stub
t1-0/2/1.0            PtToPt   0.0.0.0         0.0.0.0         0.0.0.0 0
```

```
Type P2P, Address 0.0.0.0, Mask 0.0.0.0, MTU 1500, Cost 2604
  Adj count 0
Hello 10, Dead 40, ReXmit 5, Not Stub
Auth type: MD5, Active key ID 3, Start time 2002 Nov 19 10:00:00 PST
IPsec SA Name: sa
```

show ospf3 interface detail

```
user@host> show ospf3 interface so-0/0/3.0 detail
Interface              State     Area           DR-ID         BDR-ID     Nbrs
so-0/0/3.0             PtToPt   0.0.0.0        0.0.0.0       0.0.0.0       1
Address fe80::2a0:a5ff:fe28:1dfc, Prefix-length 64
OSPF3-Intf-index 1, Type P2P, MTU 4470, Cost 12, Adj-count 1
Hello 10, Dead 40, ReXmit 5, Not Stub
```

show ospf interface detail (When Multiarea Adjacency Is Configured)

```
user@host> show ospf interface detail
regress@router> show ospf interface detail
Interface          State   Area            DR ID           BDR ID          Nbrs
lo0.0              DR      0.0.0.0         10.255.245.2    0.0.0.0            0

  Type: LAN, Address: 127.0.0.1, Mask: 255.255.255.255, MTU: 65535, Cost: 0
  DR addr: 127.0.0.1, Adj count: 0, Priority: 128
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Topology default (ID 0) -> Cost: 0
lo0.0              DR      0.0.0.0         10.255.245.2    0.0.0.0            0

  Type: LAN, Address: 10.255.245.2, Mask: 255.255.255.255, MTU: 65535, Cost: 0
  DR addr: 10.255.245.2, Adj count: 0, Priority: 128
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Topology default (ID 0) -> Cost: 0
so-0/0/0.0         PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1

  Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Topology default (ID 0) -> Cost: 1
so-0/0/0.0         PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            0

  Type: P2P, Address: 192.168.37.46, Mask: 255.255.255.254, MTU: 4470, Cost: 1
  Adj count: 0, , Passive
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Topology default (ID 0) -> Passive, Cost: 1
so-1/0/0.0         PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            1

  Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
  Adj count: 1
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Topology default (ID 0) -> Cost: 1
so-1/0/0.0         PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            0

  Type: P2P, Address: 192.168.37.54, Mask: 255.255.255.254, MTU: 4470, Cost: 1
  Adj count: 0, , Passive
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Topology default (ID 0) -> Passive, Cost: 1
so-0/0/0.0         PtToPt  1.1.1.1         0.0.0.0         0.0.0.0            1

  Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
```

```
  Adj count: 1, Secondary
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Topology default (ID 0) -> Cost: 1
so-1/0/0.0          PtToPt  1.1.1.1         0.0.0.0         0.0.0.0            1

  Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
  Adj count: 1, Secondary
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Topology default (ID 0) -> Cost: 1
so-0/0/0.0          PtToPt  2.2.2.2         0.0.0.0         0.0.0.0            1

  Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
  Adj count: 1, Secondary
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Topology default (ID 0) -> Cost: 1
so-1/0/0.0          PtToPt  2.2.2.2         0.0.0.0         0.0.0.0            1

  Type: P2P, Address: 0.0.0.0, Mask: 0.0.0.0, MTU: 4470, Cost: 1
  Adj count: 1, Secondary
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Topology default (ID 0) -> Cost: 1
```

show ospf interface area area-id

```
user@host> show ospf interface area 1.1.1.1
Interface           State   Area            DR ID           BDR ID          Nbrs
so-0/0/0.0          PtToPt  1.1.1.1         0.0.0.0         0.0.0.0            1
so-1/0/0.0          PtToPt  1.1.1.1         0.0.0.0         0.0.0.0            1
```

show ospf interface extensive (When Flooding Reduction Is Enabled)

```
user@host> show ospf interface extensive
Interface           State   Area            DR ID           BDR ID          Nbrs
fe-0/0/0.0          PtToPt  0.0.0.0         0.0.0.0         0.0.0.0            0

  Type: P2P, Address: 10.10.10.1, Mask: 255.255.255.0, MTU: 1500, Cost: 1
  Adj count: 0
  Secondary, Flood Reduction
  Hello: 10, Dead: 40, ReXmit: 5, Not Stub
  Auth type: None
  Topology default (ID 0) -> Cost: 1
```

show (ospf | ospf3) io-statistics

Syntax show (ospf | ospf3) io-statistics
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch) show (ospf | ospf3) io-statistics

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display Open Shortest Path First (OSPF) input and output statistics.

Options none—Display OSPF input and output statistics.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level view

Related Documentation

- **clear (ospf | ospf3) statistics on page 1753**

List of Sample Output **show ospf io-statistics on page 1778**

Output Fields Table 204 on page 1778 lists the output fields for the **show ospf io-statistics** command. Output fields are listed in the approximate order in which they appear.

Table 204: show (ospf | ospf3) io-statistics Output Fields

Field Name	Field Description
Packets read	Number of OSPF packets read since the last time the routing protocol was started.
average per run	Total number of packets divided by the total number of times the OSPF read operation is scheduled to run.
max run	Maximum number of packets for a given run among all scheduled runs.
Receive errors	Number of faulty packets received with errors.

show ospf io-statistics

```
user@host> show ospf io-statistics

Packets read: 7361, average per run: 1.00, max run: 1
Receive errors:
  None
```

show (ospf | ospf3) log

Syntax

show (ospf | osfp3) log
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>
<topology *topology-name*>

Syntax (J-EX Series Switch)

show (ospf | osfp3) log
<instance *instance-name*>
<topology *topology-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display the entries in the Open Shortest Path First (OSPF) log of SPF calculations.

Options

none—Display entries in the OSPF log of SPF calculations for all routing instances.

instance *instance-name*—(Optional) Display entries for the specified routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

topology *topology-name*—(Optional) Display entries for the specified topology.

realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)—(OSPFv3 only) (Optional) Display entries for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

Required Privilege Level

view

List of Sample Output

show ospf log on page 1780
show ospf log topology voice on page 1780

Output Fields

Table 205 on page 1779 lists the output fields for the **show (ospf | ospf3) log** command. Output fields are listed in the approximate order in which they appear.

Table 205: show (ospf | ospf3) log Output Fields

Field Name	Field Description
When	Time, in weeks (**w**) and days (**d**), since the SPF calculation was made.
Type	Type of calculation: **Cleanup**, **External**, **Interarea**, **NSSA**, **Redist**, **SPF**, **Stub**, **Total**, or **Virtuallink**.
Elapsed	Amount of time, in seconds, that elapsed during the operation, or the time required to complete the SPF calculation. The start time is the time displayed in the **When** field.

show ospf log

```
user@host> show ospf log
When            Type            Elapsed

1w4d 17:25:58   Stub            0.000017
1w4d 17:25:58   SPF             0.000070
1w4d 17:25:58   Stub            0.000019
1w4d 17:25:58   Interarea       0.000054
1w4d 17:25:58   External        0.000005
1w4d 17:25:58   Cleanup         0.000203
1w4d 17:25:58     Total         0.000537
1w4d 17:24:48   SPF             0.000125
1w4d 17:24:48   Stub            0.000017
1w4d 17:24:48   SPF             0.000100
1w4d 17:24:48   Stub            0.000016
1w4d 17:24:48   Interarea       0.000056
1w4d 17:24:48   External        0.000005
1w4d 17:24:48   Cleanup         0.000238
1w4d 17:24:48     Total         0.000600
...
```

show ospf log topology voice

```
user@host> show ospf log topology voice
Topology voice SPF log:

   Last instance of each event type
When            Type            Elapsed
00:06:11        SPF             0.000116
00:06:11        Stub            0.000114
00:06:11        Interarea       0.000126
00:06:11        External        0.000067
00:06:11        NSSA            0.000037
00:06:11        Cleanup         0.000186

   Maximum length of each event type
When            Type            Elapsed
00:13:43        SPF             0.000140
00:13:33        Stub            0.000116
00:13:43        Interarea       0.000128
00:13:33        External        0.000075
00:13:38        NSSA            0.000039
00:13:53        Cleanup         0.000657

   Last 100 events
When            Type            Elapsed

00:13:53        SPF             0.000090
00:13:53        Stub            0.000041
00:13:53        Interarea       0.000123
00:13:53        External        0.000040
00:13:53        NSSA            0.000038
00:13:53        Cleanup         0.000657
00:13:53          Total         0.001252
.
.
00:06:11        SPF             0.000116
00:06:11        Stub            0.000114
00:06:11        Interarea       0.000126
00:06:11        External        0.000067
00:06:11        NSSA            0.000037
```

```
00:06:11        Cleanup          0.000186
00:06:11          Total          0.000818
```

show (ospf | ospf3) neighbor

Syntax

show (ospf | ospf3) neighbor
<brief | detail | extensive>
<area *area-id*>
<instance (all | *instance-name*)>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>
<*neighbor*>
<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>

Syntax (J-EX Series Switch)

show (ospf | ospf3) neighbor
<brief | detail | extensive>
<area *area-id*>
<instance (all | *instance-name*)>
<interface *interface-name*>
<*neighbor*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display information about Open Shortest Path First (OSPF) neighbors.

Options

none—Display standard information about all OSPF neighbors for all routing instances.

brief | detail | extensive—(Optional) Display the specified level of output.

area *area-id*—(Optional) Display information about the OSPF neighbors for the specified area.

instance (all | *instance-name*—(Optional) Display all OSPF interfaces for all routing instances or under the named routing instance.

interface *interface-name*—(Optional) Display information about OSPF neighbors for the specified logical interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

neighbor—(Optional) Display information about the specified OSPF neighbor.

realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)—(Optional) (OSPFv3 only) Display information about the OSPF neighbors for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

Required Privilege Level

view

Related Documentation

- **clear (ospf | ospf3) neighbor on page 1752**

List of Sample Output

show ospf neighbor brief on page 1784

show ospf neighbor detail on page 1784
show ospf neighbor extensive on page 1785
show ospf3 neighbor detail on page 1786
show ospf neighbor area area-id on page 1786
show ospf neighbor interface interface-name on page 1786
show ospf3 neighbor instance all (OSPFv3 Multiple Family Address Support Enabled) on page 1786

Output Fields Table 206 on page 1783 lists the output fields for the **show (ospf | ospf3) neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 206: show (ospf | ospf3) neighbor Output Fields

Field Name	Field Description	Level of Output
Address	Address of the neighbor.	All levels
Interface	Interface through which the neighbor is reachable.	All levels
State	State of the neighbor: • **Attempt**—Valid only for neighbors attached to nonbroadcast networks. It indicates that no recent information has been received from the neighbor, but that a more concerted effort must be made to contact the neighbor. • **Down**—Initial state of a neighbor conversation. It indicates that no recent information has been received from the neighbor. Hello packets might continue to be sent to neighbors in the **Down** state, although at a reduced frequency. • **Exchange**—Routing device is describing its entire link-state database by sending database description packets to the neighbor. Each packet has a sequence number and is explicitly acknowledged. • **ExStart**—First step in creating an adjacency between the two neighboring routing devices. The goal of this step is to determine which routing device is the master, and to determine the initial sequence number. • **Full**—Neighboring routing devices are fully adjacent. These adjacencies appear in router link and network link advertisements. • **Init**—A Hello packet has recently been sent by the neighbor. However, bidirectional communication has not yet been established with the neighbor. This state may occur, for example, because the routing device itself did not appear in the neighbor's hello packet. • **Loading**—Link-state request packets are sent to the neighbor to acquire more recent advertisements that have been discovered (but not yet received) in the **Exchange** state. • **2Way**—Communication between the two routing devices is bidirectional. This state has been ensured by the operation of the Hello Protocol. This is the most advanced state short of beginning adjacency establishment. The (backup) designated router is selected from the set of neighbors in state **2Way** or greater.	All levels
ID	Router ID of the neighbor.	All levels
Pri	Priority of the neighbor to become the designated router.	All levels
Dead	Number of seconds until the neighbor becomes unreachable.	All levels

Table 206: show (ospf | ospf3) neighbor Output Fields *(continued)*

Field Name	Field Description	Level of Output
Link state acknowledgment list	Number of link-state acknowledgments received.	**extensive**
Link state retransmission list	Total number of link-state advertisements retransmitted. For **extensive** output only, the following information is also displayed: • **Type**—Type of link advertisement: **ASBR**, **Sum**, **Extern**, **Network**, **NSSA**, **OpaqArea**, **Router**, or **Summary**. • **LSA ID**—LSA identifier included in the advertisement. An asterisk preceding the identifier marks database entries that originated from the local routing device. • **Adv rtr**—Address of the routing device that sent the advertisement. • **Seq**—Link sequence number of the advertisement.	**detail extensive**
Neighbor-address	(OSPFv3 only) If the neighbor uses virtual links, the **Neighbor-address** is the site-local, local, or global address. If the neighbor uses a physical interface, the **Neighbor-address** is an IPv6 link-local address.	**detail extensive**
area	Area that the neighbor is in.	**detail extensive**
OSPF3-Intf-Index	(OSPFv3 only) Displays the OSPFv3 interface index.	**detail extensive**
opt	Option bits received in the hello packets from the neighbor.	**detail extensive**
DR or **DR-ID**	Address of the designated router.	**detail extensive**
BDR or **BDR-ID**	Address of the backup designated router.	**detail extensive**
Up	Length of time since the neighbor came up.	**detail extensive**
adjacent	Length of time since the adjacency with the neighbor was established.	**detail extensive**

show ospf neighbor brief

```
user@host> show ospf neighbor brief
  Address          Intf              State     ID               Pri  Dead
192.168.254.225  fxp3.0            2Way      10.250.240.32    128   36
192.168.254.230  fxp3.0            Full      10.250.240.8     128   38
192.168.254.229  fxp3.0            Full      10.250.240.35    128   33
10.1.1.129       fxp2.0            Full      10.250.240.12    128   37
10.1.1.131       fxp2.0            Full      10.250.240.11    128   38
10.1.2.1         fxp1.0            Full      10.250.240.9     128   32
10.1.2.81        fxp0.0            Full      10.250.240.10    128   33
```

show ospf neighbor detail

```
user@host> show ospf neighbor detail
  Address          Interface             State     ID               Pri  Dead
10.5.1.2         ge-1/2/0.1            Full      10.5.1.2         128   37
  area 0.0.0.1, opt 0x42, DR 10.5.1.2, BDR 10.5.1.1
  Up 06:09:28, adjacent 05:17:36
    Link state acknowledgment list:  3 entries
```

```
    Link state retransmission list:  9 entries

10.5.10.2        ge-1/2/0.10            ExStart   10.5.1.38        128   34
  area 0.0.0.1, opt 0x42, DR 10.5.10.2, BDR 10.5.10.1
  Up 06:09:28
    master, seq 0xac1530f8, rexmit DBD in 3 sec
    rexmit LSREQ in 0 sec
10.5.11.2        ge-1/2/0.11            Full      10.5.1.42        128   38
  area 0.0.0.1, opt 0x42, DR 10.5.11.2, BDR 10.5.11.1
  Up 06:09:28, adjacent 05:26:46
    Link state retransmission list:  1 entries

10.5.12.2        ge-1/2/0.12            ExStart   10.5.1.46        128   33
  area 0.0.0.1, opt 0x42, DR 10.5.12.2, BDR 10.5.12.1
  Up 06:09:28
    master, seq 0xac188a68, rexmit DBD in 2 sec
    rexmit LSREQ in 0 sec
```

show ospf neighbor extensive

```
user@host> show ospf neighbor extensive
  Address          Interface              State      ID               Pri  Dead
10.5.1.2         ge-1/2/0.1             Full      10.5.1.2         128   33
  area 0.0.0.1, opt 0x42, DR 10.5.1.2, BDR 10.5.1.1
  Up 06:09:42, adjacent 05:17:50
    Link state retransmission list:

     Type       LSA ID           Adv rtr          Seq

     Summary  10.8.56.0        172.25.27.82     0x8000004d

     Router   10.5.1.94        10.5.1.94        0x8000005c

     Network  10.5.24.2        10.5.1.94        0x80000036

     Summary  10.8.57.0        172.25.27.82     0x80000024

     Extern   1.10.90.0        10.8.1.2         0x80000041

     Extern   1.4.109.0        10.6.1.2         0x80000041

     Router   10.5.1.190       10.5.1.190       0x8000005f

     Network  10.5.48.2        10.5.1.190       0x8000003d

     Summary  10.8.58.0        172.25.27.82     0x8000004d

     Extern   1.10.91.0        10.8.1.2         0x80000041

     Extern   1.4.110.0        10.6.1.2         0x80000041

     Router   10.5.1.18        10.5.1.18        0x8000005f

     Network  10.5.5.2         10.5.1.18        0x80000033

     Summary  10.8.59.0        172.25.27.82     0x8000003a

     Summary  10.8.62.0        172.25.27.82     0x80000025

10.5.10.2        ge-1/2/0.10            ExStart   10.5.1.38        128   38
  area 0.0.0.1, opt 0x42, DR 10.5.10.2, BDR 10.5.10.1
  Up 06:09:42
    master, seq 0xac1530f8, rexmit DBD in 2 sec
```

```
    rexmit LSREQ in 0 sec
10.5.11.2        ge-1/2/0.11            Full      10.5.1.42        128   33
  area 0.0.0.1, opt 0x42, DR 10.5.11.2, BDR 10.5.11.1
  Up 06:09:42, adjacent 05:27:00
    Link state retransmission list:

      Type      LSA ID           Adv rtr          Seq

     Summary  10.8.58.0        172.25.27.82     0x8000004d

     Extern   1.10.91.0        10.8.1.2         0x80000041

     Extern   1.1.247.0        10.5.1.2         0x8000003f

     Extern   1.4.110.0        10.6.1.2         0x80000041

     Router   10.5.1.18        10.5.1.18        0x8000005f

     Network  10.5.5.2         10.5.1.18        0x80000033

     Summary  10.8.59.0        172.25.27.82     0x8000003a
```

show ospf3 neighbor detail

```
user@host> show ospf3 neighbor detail
ID               Interface              State     Pri   Dead
10.255.71.13     fe-0/0/2.0             Full      128   30
  Neighbor-address fe80::290:69ff:fe9b:e002
  area 0.0.0.0, opt 0x13, OSPF3-Intf-Index 2
  DR-ID 10.255.71.13, BDR-ID 10.255.71.12
  Up 02:51:43, adjacent 02:51:43
```

show ospf neighbor area area-id

```
user@host >show ospf neighbor area 1.1.1.1
Address          Interface              State     ID               Pri  Dead
192.168.37.47    so-0/0/0.0             Full      10.255.245.4     128    33
  Area 1.1.1.1
192.168.37.55    so-1/0/0.0             Full      10.255.245.5     128    37
  Area 1.1.1.1
```

show ospf neighbor interface interface-name

```
user@host >show ospf neighbor interface so-0/0/0.0
Address          Interface              State     ID               Pri  Dead
192.168.37.47    so-0/0/0.0             Full      10.255.245.4     128    37
  Area 0.0.0.0
192.168.37.47    so-0/0/0.0             Full      10.255.245.4     128    33
  Area 1.1.1.1
192.168.37.47    so-0/0/0.0             Full      10.255.245.4     128    32
  Area 2.2.2.2
```

show ospf3 neighbor instance all (OSPFv3 Multiple Family Address Support Enabled)

```
user @host > show ospf3 neighbor instance all
Instance: ina
    Realm: ipv6-unicast
    ID              Interface             State     Pri   Dead
   100.1.1.1       fe-0/0/2.0            Full      128     37
       Neighbor-address fe80::217:cb00:c87c:8c03
   Instance: inb
   Realm: ipv4-unicast
   ID              Interface             State     Pri   Dead
  100.1.2.1       fe-0/0/2.1            Full      128     33
      Neighbor-address fe80::217:cb00:c97c:8c03
```

show (ospf | ospf3) overview

Syntax

show (ospf | ospf3) overview
<brief | extensive>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>

Syntax (J-EX Series Switch)

show (ospf | ospf3) overview
<brief | extensive>
<instance *instance-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.
Database protection introduced in Junos OS Release 10.2.

Description

Display Open Shortest Path First (OSPF) overview information.

Options

none—Display standard information about all OSPF neighbors for all routing instances.

brief | extensive—(Optional) Display the specified level of output.

instance *instance-name*—(Optional) Display all OSPF interfaces under the named routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)—(Optional) (OSPFv3 only) Display information about the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

Required Privilege Level

view

List of Sample Output

show ospf overview on page 1789
show ospf overview (with Database Protection) on page 1789
show ospf3 overview (with database protection) on page 1790
show ospf overview extensive on page 1790

Output Fields

Table 207 on page 1787 lists the output fields for the **show ospf overview** command. Output fields are listed in the approximate order in which they appear.

Table 207: show ospf overview Output Fields

Field name	Field Description	Level of Output
Instance	OSPF routing instance.	All levels
Router ID	Router ID of the routing device.	All levels
Route table index	Route table index.	All levels

Table 207: show ospf overview Output Fields *(continued)*

Field name	Field Description	Level of Output
Configured overload	Overload capability is enabled. If the overload timer is also configured, display the time that remains before it is set to expire. This field is not displayed after the timer expires.	All levels
Full SPF runs	Number of complete Shortest Path First calculations.	All levels
SPF delay	Delay before performing consecutive Shortest Path First calculations.	All levels
SPF holddown	Delay before performing additional Shortest Path First (SPF) calculations after the maximum number of consecutive SPF calculations is reached.	All levels
SPF rapid runs	Maximum number of Shortest Path First calculations that can be performed in succession before the holddown timer begins.	All levels
LSA refresh time	Refresh period for link-state advertisement (in minutes).	All levels
Database protection state	Current state of database protection.	All levels
Warning threshold	Threshold at which a warning message is logged (percentage of maximum LSA count).	All levels
Non self-generated LSAs	Number of LSAs whose router ID is not equal to the local router ID: **Current**, **Warning** (threshold), and **Allowed**.	All levels
Ignore time	How long the database has been in the ignore state.	All levels
Reset time	How long the database must stay out of the ignore or isolated state before it returns to normal operations.	All levels
Ignore count	Number of times the database has been in the ignore state: **Current** and **Allowed**.	All levels
Restart	Graceful restart capability: **enabled** or **disabled**.	All levels
Restart duration	Time period for complete reacquisition of OSPF neighbors.	All levels
Restart grace period	Time period for which the neighbors should consider the restarting routing device as part of the topology.	All levels
Helper mode	Graceful restart helper capability: **enabled** or **disabled**.	All levels
Trace options	OSPF-specific trace options.	**extensive**
Trace file	Name of the file to receive the output of the tracing operation.	**extensive**
Area	Area number. Area 0.0.0.0 is the backbone area.	All levels
Stub type	Stub type of area: **Normal Stub**, **Not Stub**, or **Not so Stubby Stub**.	All levels

Table 207: show ospf overview Output Fields *(continued)*

Field name	Field Description	Level of Output
Authentication Type	Type of authentication: **None**, **Password**, or **MD5**.	All levels
Area border routers	Number of area border routers.	All levels
Neighbors	Number of autonomous system boundary routers.	All levels

show ospf overview

```
user@host> show ospf overview
Instance: master
    Router ID: 10.255.245.6
    Route table index: 0
   Configured overload, expires in 118 seconds
    LSA refresh time: 50 minutes
  Restart: Enabled
     Restart duration: 20 sec
     Restart grace period: 40 sec
     Helper mode: enabled
  Area: 0.0.0.0
      Stub type: Not Stub
      Authentication Type: None
      Area border routers: 0, AS boundary routers: 0
      Neighbors
        Up (in full state): 0
    Topology: default (ID 0)
      Prefix export count: 0
      Full SPF runs: 1
      SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
```

show ospf overview (with Database Protection)

```
user@host> show ospf overview
Instance: master
  Router ID: 10.255.112.218
  Route table index: 0
  LSA refresh time: 50 minutes
  Traffic engineering
  Restart: Enabled
    Restart duration: 180 sec
    Restart grace period: 210 sec
        Helper mode: Enabled
  Database protection state: Normal
    Warning threshold: 70 percent
    Non self-generated LSAs: Current 582, Warning 700, Allowed 1000
    Ignore time: 30, Reset time: 60
    Ignore count: Current 0, Allowed 1
  Area: 0.0.0.0
    Stub type: Not Stub
    Authentication Type: None
    Area border routers: 0, AS boundary routers: 0
    Neighbors
      Up (in full state): 160
  Topology: default (ID 0)
    Prefix export count: 0
    Full SPF runs: 70
```

```
    SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
    Backup SPF: Not Needed
```

show ospf3 overview (with database protection)

```
user@host> show ospf3 overview
Instance: master
  Router ID: 10.255.112.128
  Route table index: 0
  LSA refresh time: 50 minutes
  Database protection state: Normal
    Warning threshold: 80 percent
    Non self-generated LSAs: Current 3, Warning 8, Allowed 10
    Ignore time: 30, Reset time: 60
    Ignore count: Current 0, Allowed 2
  Area: 0.0.0.0
    Stub type: Not Stub
    Area border routers: 0, AS boundary routers: 0
    Neighbors
      Up (in full state): 1
  Topology: default (ID 0)
    Prefix export count: 0
    Full SPF runs: 7
    SPF delay: 0.200000 sec, SPF holddown: 5 sec, SPF rapid runs: 3
    Backup SPF: Not Needed
```

show ospf overview extensive

```
user@host> show ospf overview extensive
Instance: master
  Router ID: 1.1.1.103
  Route table index: 0
  Full SPF runs: 13, SPF delay: 0.200000 sec
  LSA refresh time: 50 minutes
  Restart: Disabled
  Trace options: lsa
  Trace file: /var/log/ospf size 131072 files 10
  Area: 0.0.0.0
    Stub type: Not Stub
    Authentication Type: None
    Area border routers: 0, AS boundary routers: 0
    Neighbors
      Up (in full state): 1
```

show (ospf | ospf3) route

Syntax

show (ospf | ospf3) route
<brief | detail | extensive>
<abr | asbr | extern | inter | intra>
<instance *instance-name*
<logical-system (all | *logical-system-name*)>
<network>
<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>
<router>
<topology *topology-name*>
<transit>

Syntax (J-EX Series Switch)

show (ospf | ospf3) route
<brief | detail | extensive>
<abr | asbr | extern | inter | intra>
<instance *instance-name*
<network>
<router>
<topology *topology-name*>
<transit>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display the entries in the Open Shortest Path First (OSPF) routing table.

Options

none—Display standard information about all entries in the OSPF routing table for all routing instances and all topologies.

brief | detail | extensive—(Optional) Display the specified level of output.

abr—(Optional) Display routes to area border routers.

asbr—(Optional) Display routes to autonomous system border routers.

extern—(Optional) Display external routes.

inter—(Optional) Display interarea routes.

intra—(Optional) Display intra-area routes.

instance *instance-name*—(Optional) Display entries for the specified routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

network—(Optional) Display routes to networks.

realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)—(OSPFv3 only) (Optional) Display entries in the routing table for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

router—(Optional) Display routes to all routers.

topology *topology-name*—(OSPF only) (Optional) Display routes for a particular topology.

transit—(Optional) (OSPFv3 only) Display OSPFv3 routes to pseudonodes.

Required Privilege Level view

List of Sample Output
show ospf route on page 1793
show ospf route detail on page 1793
show ospf3 route on page 1794
show ospf3 route detail on page 1794
show ospf route topology voice on page 1795

Output Fields Table 208 on page 1792 list the output fields for the **show (ospf | ospf3) route** command. Output fields are listed in the approximate order in which they appear.

Table 208: show (ospf | ospf3) route Output Fields

Field Name	Field Description	Output Level
Topology	Name of the topology.	All levels
Prefix	Destination of the route.	All levels
Path type	How the route was learned: • **Inter**—Interarea route • **Ext1**—External type 1 route • **Ext2**—External type 2 route • **Intra**—Intra-area route	All levels
Route type	The type of routing device from which the route was learned: • **AS BR**—Route to AS border router • **Area BR**—Route to area border router • **Area/AS BR**—Route to router that is both an **Area BR** and **AS BR**. • **Network**—Network router. • **Router**—Route to a router that is neither an **Area BR** nor an **AS BR**. • **Transit**—(OSPFv3 only) Route to a pseudonode representing a transit network, LAN, or nonbroadcast multiaccess (NBMA) link. • **Discard**—Route to a summary discard.	All levels
NH Type	Next-hop type: **LSP** or **IP**.	All levels
Metric	Route's metric value.	All levels
NH-interface	(OSPFv3 only) Interface through which the route's next hop is reachable.	All levels
NH-addr	(OSPFv3 only) IPv6 address of the next hop.	All levels

Table 208: show (ospf | ospf3) route Output Fields *(continued)*

Field Name	Field Description	Output Level
NextHop Interface	(OSPFv2 only) Interface through which the route's next hop is reachable.	All levels
Nexthop addr/label	(OSPFv2 only) If the **NH Type** is **IP**, then it is the address of the next hop. If the **NH Type** is **LSP**, then it is the name of the label-switched path.	All levels
Area	Area ID of the route.	**detail**
Origin	Router from which the route was learned.	**detail**
Type 7	Route was learned through a not-so-stubby area (NSSA) link-state advertisement (LSA).	**detail**
P-bit	Route was learned through NSSA LSA and the propagate bit was set.	**detail**
Fwd NZ	Forwarding address is nonzero. **Fwd NZ** is only displayed if the route is learned through an NSSA LSA.	**detail**
optional-capability	Optional capabilities propagated in the router LSA. This field is in the output for intraarea router routes only (when **Route Type** is **Area BR**, **AS BR**, **Area/AS BR**, or **Router**), not for interarea router routes or network routes. Three bits in this field are defined as follows: • **0x4 (V)**—Routing device is at the end of a virtual active link. • **0x2 (E)**—Routing device is an autonomous system boundary router. • **0x1 (B)**—Routing device is an area border router.	**detail**
priority	The priority assigned to the prefix: • **high** • **medium** • **low** NOTE: The **priority** field applies only to routes of type **Network**.	**detail**

show ospf route

```
user@host> show ospf route
Prefix                 Path   Route       NH    Metric  NextHop       Nexthop
                              Type   Type         Type         Interface
addr/label
10.255.71.12           Intra  Router       IP   1       fe-0/0/2.0    192.16.22.86

10.255.71.13/32        Intra  Network      IP   0       lo0.0
192.168.222.84/30      Intra  Network      LSP  1       fe-0/0/2.0    lsp-ab
```

show ospf route detail

```
user@host> show ospf route detail
Topology default Route Table:

Prefix             Path   Route       NH    Metric  NextHop       Nexthop
                    Type   Type         Type         Interface     addr/
label
10.255.14.174      Inter  AS BR       IP      210  t1-3/0/1.0
   area 0.0.0.2, origin 10.255.14.185
```

```
10.255.14.178      Intra  Router      IP      200  t3-3/1/3.0
   area 0.0.0.2, origin 10.255.14.178, optional-capability 0x0
10.210.1.0/30      Intra  Network     IP       10  t3-3/1/2.0
   area 0.0.0.2, origin 10.255.14.172, priority medium
100.1.1.1/32       Inter  Network     IP      210  t1-3/0/1.0
   area 0.0.0.2, origin 10.255.14.185, priority low
112.3.1.0/24       Ext2   Network     IP        0  t1-3/0/1.0
   area 0.0.0.0, origin 10.255.14.174, priority high
200.3.3.0/30       Inter  Network     IP      220  t1-3/0/1.0
   area 0.0.0.2, origin 10.255.14.185, priority high
```

show ospf3 route

```
user@host> show ospf3 route
Prefix                                  Path   Route       NH   Metric
                                        type   type        type
10.255.71.13                            Intra  Router      IP   1
  NH-interface fe-0/0/2.0, NH-addr fe80::290:69ff:fe9b:e002
10.255.71.13;0.0.0.2              Prefix             Path   Route       NH
Metric  NextHop       Nexthop
                    Type   Type        Type         Interface     addr/label
10.255.245.1       Intra  Router      IP       40  fxp1.1        192.168.36.17
  area 0.0.0.0, origin 10.255.245.1 optional-capability 0x0,
10.255.245.3       Intra  AS BR       IP        1  fxp2.3        192.168.36.34
  area 0.0.0.0, origin 10.255.245.3 optional-capability 0x0,
10.255.245.1/32    Intra  Network     IP       40  fxp1.1        192.168.36.17
  area 0.0.0.0, origin 10.255.245.1, priority high
10.255.245.2/32    Intra  Network     IP        0  lo0.0
  area 0.0.0.0, origin 10.255.245.2, priority medium
10.255.245.3/32    Intra  Network     IP        1  fxp2.3        192.168.36.34
  area 0.0.0.0, origin 10.255.245.3, priority low

          Intra  Transit     IP   1
  NH-interface fe-0/0/2.0
192::168:222:84/126                     Intra  Network     IP   1
  NH-interface fe-0/0/2.0
abcd::71:12/128                         Intra  Network     IP   0
  NH-interface lo0.0
abcd::71:13/128                         Intra  Network     LSP   1
  NH-interface fe-0/0/2.0, NH-addr lsp-cd
```

show ospf3 route detail

```
user@host>  show ospf3 route detail
Prefix                                  Path   Route       NH
Metric
                                         type   type        type
10.255.14.174                           Intra  Area/AS BR  IP   110
   NH-interface so-1/2/2.0
   Area 0.0.0.0, Origin 10.255.14.174, Optional-capability 0x3
10.255.14.178                           Intra  Router      IP   200
   NH-interface t3-3/1/3.0
   Area 0.0.0.0, Origin 10.255.14.178, Optional-capability 0x0
10.255.14.185;0.0.0.2                   Intra  Transit     IP   200
   NH-interface t1-3/0/1.0
   NH-interface so-1/2/2.0
   Area 0.0.0.0, Origin 10.255.14.185
1000:1:1::1/128                         Inter  Network     IP   110
   NH-interface so-1/2/2.0
   Area 0.0.0.0, Origin 10.255.14.174, Priority low
1001:2:1::/48                           Ext1   Network     IP   110
   NH-interface so-1/2/2.0
   Area 0.0.0.0, Origin 10.255.14.174, Fwd NZ, Priority medium
1002:1:7::/48                           Ext2   Network     IP   0
   NH-interface so-1/2/2.0
```

```
   Area 0.0.0.0, Origin 10.255.14.174, Fwd NZ, Priority low
1002:3:4::/48                                Ext2   Network     IP   0
   NH-interface so-1/2/2.0
   Area 0.0.0.0, Origin 10.255.14.174, Fwd NZ, Priority high
abcd::10:255:14:172/128                      Intra  Network     IP   0
   NH-interface lo0.0
   Area 0.0.0.0, Origin 10.255.14.172, Priority low
```

show ospf route topology voice

```
user@host show ospf route topology voice
Topology voice Route Table:
Prefix             Path   Route       NH   Metric  NextHop       Nexthop
                   Type   Type        Type         Interface     addr/label
10.255.8.2         Intra  Router      IP        1  so-0/2/0.0
10.255.8.3         Intra  Router      IP        2  so-0/2/0.0
10.255.8.1/32      Intra  Network     IP        0  lo0.0
10.255.8.2/32      Intra  Network     IP        1  so-0/2/0.0
10.255.8.3/32      Intra  Network     IP        2  so-0/2/0.0
192.168.8.0/29     Intra  Network     IP        2  so-0/2/0.0
192.168.8.44/30    Intra  Network     IP        2  so-0/2/0.0
192.168.8.46/32    Intra  Network     IP        1  so-0/2/0.0
192.168.8.48/30    Intra  Network     IP        1  so-0/2/1.0
192.168.8.52/30    Intra  Network     IP        2  so-0/2/0.0
192.168.9.44/30    Intra  Network     IP        1  so-0/2/0.0
192.168.9.45/32    Intra  Network     IP        2  so-0/2/0.0
```

show (ospf | ospf3) statistics

Syntax

show (ospf | ospf3) statistics
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>

Syntax (J-EX Series Switch)

show (ospf | ospf3) statistics
<instance *instance-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display OSPF statistics.

Options

none—Display OSPF statistics for all routing instances.

instance *instance-name*—(Optional) Display all statistics for the specified routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)—(Optional) (OSPFv3 only) Display all statistics for the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family for OSPFv3 other than IPv6 unicast, which is the default.

Required Privilege Level

view

Related Documentation

- **clear (ospf | ospf3) statistics on page 1753**

List of Sample Output

show ospf statistics on page 1797

Output Fields

Table 209 on page 1796 lists the output fields for the **show (ospf | ospf3) statistics** command. Output fields are listed in the approximate order in which they appear.

Table 209: show (ospf | ospf3) statistics Output Fields

Field Name	Field Description
Packet type	Type of OSPF packet.
Total Sent/Total Received	Total number of packets sent and received.
Last 5 seconds Sent/Last 5 seconds Received	Total number of packets sent and received in the last 5 seconds.
LSAs retransmitted	Total number of link-state advertisements transmitted, and number retransmitted in the last 5 seconds.
Receive errors	Number and type of receive errors.

show ospf statistics

```
user@host> show ospf statistics
Packet type              Total                  Last 5 seconds
                      Sent      Received        Sent      Received
   Hello          505739        990495           4             5
     DbD              20            26           0             0
   LSReq               6             5           0             0
LSUpdate           27060         15319           0             0
   LSAck           10923         52470           0             0

LSAs retransmitted: 16, last 5 seconds: 0

Receive errors:
  862 no interface found
  115923 no virtual link found
```

show as-path

Syntax
show as-path
<brief | detail>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show as-path
<brief | detail>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display the distribution of autonomous system (AS) paths that the local routing device is using (usually through the routing table). Use this command to debug problems for AS paths and to understand how AS paths have been manipulated through a policy (through the **as-path-prepend** action) or through aggregation.

Options
none—Display basic information about AS paths that the local routing device is using (same as brief).

brief | detail—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show as-path on page 1799
show as-path detail on page 1800

Output Fields
Table 210 on page 1798 lists the output fields for the **show as-path** command. Output fields are listed in the approximate order in which they appear.

Table 210: show as-path Output Fields

Field Name	Field Description	Level of Output
Total AS paths	Total number of AS paths.	**brief** none
Bucket	Bucket value. This value represents a traffic classification on the interface.	All levels
Count	Path reference count.	All levels
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **?**—Incomplete; typically, the AS path was aggregated. • **Atomic**—Route is an aggregate of several route prefixes. • **Aggregator**—Routing device has summarized a range of prefixes.	All levels

Table 210: show as-path Output Fields *(continued)*

Field Name	Field Description	Level of Output
domain	Number of independent AS domains. The AS paths of an independent AS domain are not shared with the AS paths and AS path attributes of other domains, including the master routing instance domain.	**detail**
neighbor as	AS peer address.	**detail**
length	Length of the AS path.	**detail**
segments	Length of the AS segment descriptor.	**detail**
references	Path reference count.	**detail**

show as-path

```
user@host> show as-path
Total AS paths: 30382
  Bucket 0    Count: 36
    I
    14203 2914 174 31752 I
    14203 2914 701 21512 I
    14203 2914 1239 26632 I
    14203 2914 1239 29704 I
    14203 2914 4323 10248 I
    14203 2914 4766 23560 I
    14203 2914 6395 32776 I
    14203 2914 7911 11272 I
    14203 2914 12180 18440 I
    14203 2914 17408 17416 I
    14203 2914 701 702 24586 I
    14203 2914 1239 4657 9226 I
    14203 2914 1239 7132 16394 I
    14203 2914 1299 8308 34826 I
    14203 2914 3320 5603 28682 I
    14203 2914 3491 1680 33802 I
    14203 2914 3549 7908 27658 I
    14203 2914 3549 20804 30730 I
    14203 2914 7018 2687 9226 I
    14203 2914 174 9318 9318 23564 I
    14203 2914 701 3786 3786 23564 I
    14203 2914 701 4761 4795 9228 I
    14203 2914 1239 7132 5673 18444 I
    14203 2914 3491 20485 24588 24588 I
    14203 2914 5511 2200 1945 2060 I
    14203 2914 7911 14325 14325 14348 I
    14203 2914 701 4637 9230 9230 9230 I
    14203 2914 6395 14 14 14 14 I
    14203 2914 9299 6163 6163 6163 6163 9232 I
    14203 2914 3356 3356 3356 3356 3356 11955 21522 I
    14203 2914 9837 9837 9219 I Aggregator: 9219 202.27.91.253
    14203 2914 174 30209 30222 30222 30222 ?
    14203 2914 1299 5377 I (Atomic) Aggregator: 5377 193.219.192.22
    14203 2914 4323 36097 I (Atomic) Aggregator: 36097 216.69.252.254
    14203 2914 209 2516 17676 23813 I (Atomic) Aggregator: 23813 219.127.233.66
  Bucket 1    Count: 28
    14203 2914 35847 I
    14203 2914 174 19465 I
```

```
    14203 2914 174 35849 I
    14203 2914 2828 32777 I
    14203 2914 4323 14345 I
    14203 2914 4323 29705 I
    14203 2914 6395 32777 I

...
```

show as-path detail

```
user@host> show as-path detail
Total AS paths: 30410
  Bucket 0    Count: 36
    AS path: I
      domain 0, length 0, segments 0, references 54
    AS path: 14203 2914 174 31752 I
      domain 1, neighbor as: 14203, length 4, segments 1, references 2
    AS path: 14203 2914 701 21512 I
      domain 1, neighbor as: 14203, length 4, segments 1, references 2
    AS path: 14203 2914 1239 26632 I
      domain 1, neighbor as: 14203, length 4, segments 1, references 2
    AS path: 14203 2914 1239 29704 I
      domain 1, neighbor as: 14203, length 4, segments 1, references 2
    AS path: 14203 2914 4323 10248 I
      domain 1, neighbor as: 14203, length 4, segments 1, references 2
    AS path: 14203 2914 4766 23560 I
      domain 1, neighbor as: 14203, length 4, segments 1, references 2
    AS path: 14203 2914 6395 32776 I
      domain 1, neighbor as: 14203, length 4, segments 1, references 3
    AS path: 14203 2914 7911 11272 I
      domain 1, neighbor as: 14203, length 4, segments 1, references 2
    AS path: 14203 2914 12180 18440 I
      domain 1, neighbor as: 14203, length 4, segments 1, references 3
    AS path: 14203 2914 17408 17416 I
      domain 1, neighbor as: 14203, length 4, segments 1, references 3
    AS path: 14203 2914 701 702 24586 I
      domain 1, neighbor as: 14203, length 5, segments 1, references 3
    AS path: 14203 2914 1239 4657 9226 I
      domain 1, neighbor as: 14203, length 5, segments 1, references 7
    AS path: 14203 2914 1239 7132 16394 I
      domain 1, neighbor as: 14203, length 5, segments 1, references 2
    AS path: 14203 2914 1299 8308 34826 I
      domain 1, neighbor as: 14203, length 5, segments 1, references 2
    AS path: 14203 2914 3320 5603 28682 I
      domain 1, neighbor as: 14203, length 5, segments 1, references 2
    AS path: 14203 2914 3491 1680 33802 I
      domain 1, neighbor as: 14203, length 5, segments 1, references 2
    AS path: 14203 2914 3549 7908 27658 I
      domain 1, neighbor as: 14203, length 5, segments 1, references 2
    AS path: 14203 2914 3549 20804 30730 I
      domain 1, neighbor as: 14203, length 5, segments 1, references 2
    AS path: 14203 2914 7018 2687 9226 I
      domain 1, neighbor as: 14203, length 5, segments 1, references 3
    AS path: 14203 2914 174 9318 9318 23564 I
      domain 1, neighbor as: 14203, length 6, segments 1, references 2
    AS path: 14203 2914 701 3786 3786 23564 I
      domain 1, neighbor as: 14203, length 6, segments 1, references 2
    AS path: 14203 2914 701 4761 4795 9228 I
      domain 1, neighbor as: 14203, length 6, segments 1, references 14
    AS path: 14203 2914 1239 7132 5673 18444 I
      domain 1, neighbor as: 14203, length 6, segments 1, references 2
    AS path: 14203 2914 3491 20485 24588 24588 I
      domain 1, neighbor as: 14203, length 6, segments 1, references 4
```

```
    AS path: 14203 2914 5511 2200 1945 2060 I
      domain 1, neighbor as: 14203, length 6, segments 1, references 2
    AS path: 14203 2914 7911 14325 14325 14348 I
      domain 1, neighbor as: 14203, length 6, segments 1, references 2
    AS path: 14203 2914 701 4637 9230 9230 9230 I
      domain 1, neighbor as: 14203, length 7, segments 1, references 3
    AS path: 14203 2914 6395 14 14 14 14 I
      domain 1, neighbor as: 14203, length 7, segments 1, references 10
...
```

show as-path domain

Syntax
show as-path domain
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show as-path domain

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display autonomous system (AS) path domain information.

Options
none—(Optional) Display AS path domain information for all routing instances.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show as-path domain on page 1803

Output Fields
Table 211 on page 1802 lists the output fields for the **show as-path domain** command. Output fields are listed in the approximate order in which they appear

Table 211: show as-path domain Output Fields

Field Name	Field Description
Domain	Number of independent AS domains. The AS paths of an independent AS domain are not shared with the AS paths and AS path attributes of other domains, including the master routing instance domain.
Primary	Primary AS number.
References	Path reference count.
***Number* Paths**	Number of known AS paths.
Flags	Information about the AS path: • **ASLoop**—Path contains an AS loop. • **Atomic**—Path includes the ATOMIC_AGGREGATE path attribute. • **Local**—Path was created by local aggregation. • **Master**—Path was created by the master routing instance.
Local AS	AS number of the local routing device.
Loops	How many times this AS number can appear in an AS path.

show as-path domain

```
user@host> show as-path domain
Domain: 1        Primary: 10458
    References:          3 Paths:      30383
    Flags: Master
    Local AS: 10458  Loops: 1
```

show as-path summary

Syntax
show as-path summary
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show as-path summary

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display autonomous system (AS) path summary information.

Options
none—(Optional) Display AS path summary information for all routing instances.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show as-path summary on page 1804

Output Fields
Table 212 on page 1804 lists the output fields for the **show as-path summary** command. Output fields are listed in the approximate order in which they appear.

Table 212: show as-path summary Output Fields

Field Name	Field Description
AS Path	AS path number.
Buckets	Bucket value. This value represents a traffic classification on the interface.
Max	Maximum limit for the number of AS numbers.
Min	Minimum limit for the number of AS numbers.
Avg	Average number of AS numbers.
Std deviation	Standard deviation for the number of AS numbers.

show as-path summary

```
user@host> show as-path summary
AS Paths  Buckets     Max     Min     Avg   Std deviation
   30425     1024      95      12      29        6.481419
```

show bgp bmp

Syntax **show bgp bmp**

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display information about the BGP Monitoring Protocol (BMP).

Options This command has no options.

Required Privilege Level view

List of Sample Output **show bgp bmp on page 1805**

Output Fields Table 213 on page 1805 lists the output fields for the **show bgp bmp** command. Output fields are listed in the approximate order in which they appear.

Table 213: show bgp bmp Output Fields

Field Name	Field Description
BMP station address/port:	IP address and port number of monitoring station to which BGP Monitoring Protocol (BMP) statistics are sent.
BMP session state	Status of the BMP session: **UP** or **DOWN**.
Memory consumed by BMP	Memory used by the active BMP session.
Statistics timeout	Amount of time, in seconds, between transmissions of BMP data to the monitoring station.
Memory limit	Threshold, in bytes, at which the routing device stops collecting BMP data if it is exceeded.
Memory-connect retry timeout	Amount of time, in seconds, after which the routing device attempts to resume a BMP session that was ended after the configured memory threshold was exceeded.

show bgp bmp

```
user@host> show bgp bmp
  BMP station address/port: 172.24.24.157+5454
  BMP session state: DOWN
  Memory consumed by BMP: 0
  Statistics timeout: 15
  Memory limit: 10485760
  Memory connect retry timeout: 600
```

show bgp group

Syntax

show bgp group
<brief | detail | summary>
<*group-name*>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<rtf>

Syntax (J-EX Series Switch)

show bgp group
<brief | detail | summary>
<*group-name*>
<instance *instance-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display information about the configured Border Gateway Protocol (BGP) groups.

Options

none—Display group information about all BGP groups.

brief | detail | summary—(Optional) Display the specified level of output.

group-name—(Optional) Display group information for the specified group.

instance *instance-name*—(Optional) Display information about a particular BGP peer in the specified instance. The instance name can be master for the main instance, or any valid configured instance name or its prefix.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

rtf—(Optional) Display BGP group route targeting information.

Required Privilege Level

view

List of Sample Output

show bgp group on page 1809
show bgp group on page 1810
show bgp group brief on page 1810
show bgp group detail on page 1810
show bgp group rtf detail on page 1811
show bgp group summary on page 1811
show bgp group summary on page 1811

Output Fields

Table 214 on page 1806 describes the output fields for the **show bgp group** command. Output fields are listed in the approximate order in which they appear.

Table 214: show bgp group Output Fields

Field Name	Field Description	Level of Output
Group type or **Group**	Type of BGP group: **Internal** or **External**.	All levels

Table 214: show bgp group Output Fields *(continued)*

Field Name	Field Description	Level of Output
AS	AS number of the peer. For internal BGP (IBGP), this number is the same as **Local AS**.	**brief detail** none
Local AS	AS number of the local routing device.	**brief detail** none
Name	Name of a specific BGP group.	**brief detail** none
Flags	Flags associated with the BGP group. This field is used by Dell Support (see “Requesting Technical Support” on page lxxi).	**brief detail** none
Export	Export policies configured for the BGP group with the **export** statement.	**brief detail** none
MED tracks IGP metric update delay	Time interval, in seconds, that updates to multiple exit discriminator (MED) are delayed. Also displays the time remaining before the interval is set to expire	All
Total peers	Total number of peers in the group.	**brief detail** none
Established	Number of peers in the group that are in the established state.	All levels
Active/Received/Accepted/Damped	Multipurpose field that displays information about BGP peer sessions. The field’s contents depend upon whether a session is established and whether an established session was established in the main routing device or in a routing instance. • If a peer is not established, the field shows the state of the peer session: **Active**, **Connect**, or **Idle**. • If a BGP session is established in the main routing device, the field shows the number of active, received, accepted, and damped routes that are received from a neighbor and appear in the **inet.0** (main) and **inet.2** (multicast) routing tables. For example, **8/10/10/2 and 2/4/4/0** indicate the following: ◦ 8 active routes, 10 received routes, 10 accepted routes, and 2 damped routes from a BGP peer appear in the **inet.0** routing table. ◦ 2 active routes, 4 received routes, 4 accepted routes, and no damped routes from a BGP peer appear in the **inet.2** routing table.	**summary**
ip-addresses	List of peers who are members of the group. The address is followed by the peer’s port number.	All levels
Route Queue Timer	Number of seconds until queued routes are sent. If this time has already elapsed, this field displays the number of seconds by which the updates are delayed.	**detail**

Table 214: show bgp group Output Fields *(continued)*

Field Name	Field Description	Level of Output
Route Queue	Number of prefixes that are queued up for sending to the peers in the group.	**detail**
inet.*number*	Number of active, received, accepted, and damped routes in the routing table. For example, **inet.0: 7/10/9/0** indicates the following: • 7 active routes, 10 received routes, 9 accepted routes, and no damped routes from a BGP peer appear in the **inet.0** routing table.	none
Table inet.*number*	Information about the routing table. • **Received prefixes**—Total number of prefixes from the peer, both active and inactive, that are in the routing table. • **Active prefixes**—Number of prefixes received from the peer that are active in the routing table. • **Suppressed due to damping**—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols. • **Advertised prefixes**—Number of prefixes advertised to a peer. • **Received external prefixes**—Total number of prefixes from the external BGP (EBGP) peers, both active and inactive, that are in the routing table. • **Active external prefixes**—Number of prefixes received from the EBGP peers that are active in the routing table. • **Externals suppressed**—Number of routes received from EBGP peers currently inactive because of damping or other reasons. • **Received internal prefixes**—Total number of prefixes from the IBGP peers, both active and inactive, that are in the routing table. • **Active internal prefixes**—Number of prefixes received from the IBGP peers that are active in the routing table. • **Internals suppressed**—Number of routes received from IBGP peers currently inactive because of damping or other reasons. • **RIB State**—Status of the graceful restart process for this routing table: **BGP restart is complete**, **BGP restart in progress**, **VPN restart in progress**, or **VPN restart is complete**.	**detail**
Groups	Total number of groups.	All levels
Peers	Total number of peers.	All levels
External	Total number of external peers.	All levels
Internal	Total number of internal peers.	All levels
Down peers	Total number of unavailable peers.	All levels
Flaps	Total number of flaps that occurred.	All levels
Table	Name of a routing table.	**brief**, none

Table 214: show bgp group Output Fields *(continued)*

Field Name	Field Description	Level of Output
Tot Paths	Total number of paths.	**brief**, none
Act Paths	Number of active routes.	**brief**, none
Suppressed	Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.	**brief**, none
History	Number of withdrawn routes stored locally to keep track of damping history.	**brief**, none
Damp State	Number of active routes with a figure of merit greater than zero, but lower than the threshold at which suppression occurs.	**brief**, none
Pending	Routes being processed by BGP import policy.	**brief**, none
Group	Group the peer belongs to in the BGP configuration.	**detail**
Receive mask	Mask of the received target included in the advertised route.	**detail**
Entries	Number of route entries received.	**detail**
Target	Route target that is to be passed by route-target filtering. If a route advertised from the provider edge (PE) routing device matches an entry in the route-target filter, the route is passed to the peer.	**detail**
Mask	Mask which specifies that the peer receive routes with the given route target.	**detail**

show bgp group

```
user@host> show bgp group
Group Type: Internal    AS: 21                    Local AS: 21
  Name: from_vpn04_to_other Index: 0              Flags: <>
  Holdtime: 0
  Total peers: 3        Established: 3
  10.255.14.178+179
  10.255.71.24+179
  10.255.14.182+179
  inet.0: 2/7/0

Group Type: External                              Local AS: 21
  Name: from_vpn04_to_vpn06 Index: 1              Flags: <Export Eval>
Export: [ internal-and-bgp ]
  Holdtime: 0
  Traffic Statistics Interval: 300
  Total peers: 1        Established: 1
  100.1.3.2+2910
  inet.0: 5/10/0

Groups: 2  Peers: 4    External: 1    Internal: 3    Down peers: 0   Flaps: 2
```

```
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0                17          7          0          0          0          0
```

show bgp group

```
user@host> show bgp group
Group Type: External                              Local AS: 65500
  Name: as65501peers    Index: 0                  Flags:  Export <Eval>
  Export: [ export-policy ]
  Holdtime: 0
  Total peers: 1        Established: 1
  192.168.4.222+179
  Trace options:  all
  Trace file: /var/log/bgp size 10485760 files 10
  inet.0: 7/10/9/0
  inet.2: 0/0/0/0

Groups: 1  Peers: 1    External: 1    Internal: 0    Down peers: 0   Flaps: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0                10          7          0          0          0          0
inet.2                0           0          0          0          0          0
```

show bgp group brief

The output for the **show bgp group brief** command is identical to that for the **show bgp group** command. For sample output, see **show bgp group on page 1809**.

show bgp group detail

```
user@host> show bgp group detail
Group Type: Internal    AS: 21                    Local AS: 21
  Name: from_vpn04_to_other Index: 0              Flags: <>
  Holdtime: 0
  Total peers: 3        Established: 3
  10.255.14.178+179
  10.255.71.24+179
  10.255.14.182+179
  Route Queue Timer: unset Route Queue: empty
  Table inet.0
    Active prefixes:              2
    Received prefixes:            7
    Suppressed due to damping:    0
    Advertised prefixes:          5

Group Type: External                              Local AS: 21
  Name: from_vpn04_to_vpn06 Index: 1              Flags:<Export Eval>
  Export: [ internal-and-bgp ]
  Holdtime: 0
  Traffic Statistics Interval: 300
  Total peers: 1        Established: 1
  100.1.3.2+2910
  Route Queue Timer: unset Route Queue: empty
  Table inet.0
    Active prefixes:              5
    Received prefixes:            10
    Suppressed due to damping:    0
    Advertised prefixes:          6

Groups: 2  Peers: 4    External: 1    Internal: 3    Down peers: 0   Flaps: 2
  Table inet.0
    Received prefixes:            17
    Active prefixes:              7
    Suppressed due to damping:    0
    Received external prefixes:   10
    Active external prefixes:     5
```

```
    Externals suppressed:         0
    Received internal prefixes:   7
    Active internal prefixes:     2
    Internals suppressed:         0
    RIB State: BGP restart is complete
```

show bgp group rtf detail

```
user@host> show bgp group rtf detail
Group: asbr
    Receive mask: 00000001
    Table: bgp.rtarget.0                    Flags: Filter   Entries: 4
        Target                             Mask
        109:1/64                           00000001
        109:2/64                           00000001
        701:1/64                           00000001
        10458:2/64                         00000001

Group: mesh_0
    Receive mask: 0000000e
    Table: bgp.rtarget.0                    Flags: Filter   Entries: 12
        Target                             Mask
        109:1/64                           00000002
        701:1/64                           00000002
        701:2/64                           00000002
        10458:1/64                         0000000e
        10458:2/64                         00000006
        10458:3/64                         00000006
        10458:5/64                         00000006
        10458:6/64                         00000004
        10458:7/64                         00000008
        10458:8/64                         00000008
        10458:10/64                        00000002
```

show bgp group summary

```
user@host> show bgp group summary
Group        Type       Peers     Established    Active/Received/Damped
from_vpn04_to_other Internal 3    3
  inet.0           : 2/7/0
from_vpn04_to_vpn06 External 1    1
  inet.0           : 5/10/0

Groups: 2  Peers: 4    External: 1    Internal: 3    Down peers: 0   Flaps: 2
  inet.0           : 7/17/0 External: 5/10/0 Internal: 2/7/0
```

show bgp group summary

```
user@host> show bgp group summary
Group        Type       Peers     Established    Active/Received/Accepted/Damped
as65501peers External   1         1
  Trace options:  all
  Trace file: /var/log/bgp size 10485760 files 10
  inet.0           : 7/10/9/0
  inet.2           : 0/0/0/0

Groups: 1  Peers: 1    External: 1    Internal: 0    Down peers: 0   Flaps: 0
  inet.0           : 7/10/9/0 External: 7/10/9/0 Internal: 0/0/0/0
  inet.2           : 0/0/0/0 External: 0/0/0/0 Internal: 0/0/0/0
```

show bgp neighbor

Syntax

show bgp neighbor
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<*neighbor-address*>
<orf (detail | *neighbor-address*)

Syntax (J-EX Series Switch)

show bgp neighbor
<instance *instance-name*>
<*neighbor-address*>
<orf (*neighbor-address* | detail)

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display information about Border Gateway Protocol (BGP) peers.

Options

none—Display information about all BGP peers.

instance *instance-name*—(Optional) Display information about BGP peers for only the specified routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

neighbor-address—(Optional) Display information for only the BGP peer at the specified IP address.

orf (detail | *neighbor-address*)—(Optional) Display outbound route-filtering information for all BGP peers or only for the BGP peer at the specified IP address. The default is to display brief output. Use the **detail** option to display detailed output.

Additional Information

For information about the **local-address, nlri**, **hold-time,** and **preference** statements, see the *Junos OS Routing Protocols Configuration Guide.*

Required Privilege Level

view

Related Documentation

- **clear bgp neighbor on page 1756**

List of Sample Output

show bgp neighbor (CLNS) on page 1818
show bgp neighbor (Layer 2 VPN) on page 1819
show bgp neighbor (Layer 3 VPN) on page 1820
show bgp neighbor neighbor-address on page 1821
show bgp neighbor neighbor-address on page 1822
show bgp neighbor orf neighbor-address detail on page 1823

Output Fields

Table 215 on page 1813 describes the output fields for the **show bgp neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 215: show bgp neighbor Output Fields

Field Name	Field Description
Peer	Address of the BGP neighbor. The address is followed by the neighbor's port number.
AS	AS number of the peer.
Local	Address of the local routing device. The address is followed by the peer's port number.
Type	Type of peer: **Internal** or **External**.
State	Current state of the BGP session: • **Active**—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message. • **Connect**—BGP is waiting for the transport protocol connection to be completed. • **Established**—The BGP session has been established, and the peers are exchanging update messages. • **Idle**—This is the first stage of a connection. BGP is waiting for a Start event. • **OpenConfirm**—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message. • **OpenSent**—BGP has sent an open message and is waiting to receive an open message from the peer.
Flags	Internal BGP flags: • **Aggregate Label**—BGP has aggregated a set of incoming labels (labels received from the peer) into a single forwarding label. • **CleanUp**—The peer session is being shut down. • **Delete**—This peer has been deleted. • **Idled**—This peer has been permanently idled. • **ImportEval**—At the last commit, this peer was identified as needing to reevaluate all received routes. • **Initializing**—The peer session is initializing. • **SendRtn**—Messages are being sent to the peer. • **Sync**—This peer is synchronized with the rest of the peer group. • **TryConnect**—Another attempt is being made to connect to the peer. • **Unconfigured**—This peer is not configured. • **WriteFailed**—An attempt to write to this peer failed.
Last state	Previous state of the BGP session: • **Active**—BGP is initiating a transport protocol connection in an attempt to connect to a peer. If the connection is successful, BGP sends an Open message. • **Connect**—BGP is waiting for the transport protocol connection to be completed. • **Established**—The BGP session has been established, and the peers are exchanging update messages. • **Idle**—This is the first stage of a connection. BGP is waiting for a Start event. • **OpenConfirm**—BGP has acknowledged receipt of an open message from the peer and is waiting to receive a keepalive or notification message. • **OpenSent**—BGP has sent an open message and is waiting to receive an open message from the peer.

Table 215: show bgp neighbor Output Fields *(continued)*

Field Name	Field Description
Last event	Last activity that occurred in the BGP session: • **Closed**—The BGP session closed. • **ConnectRetry**—The transport protocol connection failed, and BGP is trying again to connect. • **HoldTime**—The session ended because the hold timer expired. • **KeepAlive**—The local routing device sent a BGP keepalive message to the peer. • **Open**—The local routing device sent a BGP open message to the peer. • **OpenFail**—The local routing device did not receive an acknowledgment of a BGP open message from the peer. • **RecvKeepAlive**—The local routing device received a BGP keepalive message from the peer. • **RecvNotify**—The local routing device received a BGP notification message from the peer. • **RecvOpen**—The local routing device received a BGP open message from the peer. • **RecvUpdate**—The local routing device received a BGP update message from the peer. • **Start**—The peering session started. • **Stop**—The peering session stopped. • **TransportError**—A TCP error occurred.
Last error	Last error that occurred in the BGP session: • **Cease**—An error occurred, such as a version mismatch, that caused the session to close. • **Finite State Machine Error**—In setting up the session, BGP received a message that it did not understand. • **Hold Time Expired**—The session's hold time expired. • **Message Header Error**—The header of a BGP message was malformed. • **Open Message Error**—A BGP open message contained an error. • **None**—No errors occurred in the BGP session. • **Update Message Error**—A BGP update message contained an error.
Export	Name of the export policy that is configured on the peer.
Import	Name of the import policy that is configured on the peer.

Table 215: show bgp neighbor Output Fields *(continued)*

Field Name	Field Description
Options	Configured BGP options: • **AddressFamily**—Configured address family: **inet** or **inet-vpn**. • **AutheKeyChain**—Authentication key change is enabled. • **GracefulRestart**—Graceful restart is configured. • **HoldTime**—Hold time configured with the **hold-time** statement. The hold time is three times the interval at which keepalive messages are sent. • **Local Address**—Address configured with the **local-address** statement. • **Multihop**—Allow BGP connections to external peers that are not on a directly connected network. • **NLRI**—Configured MBGP state for the BGP group: **multicast**, **unicast**, or both if you have configured **nlri any**. • **Peer AS**—Configured peer autonomous system (AS). • **Preference**—Preference value configured with the **preference** statement. • **Refresh**—Configured to refresh automatically when the policy changes. • **Rib-group**—Configured routing table group.
Authentication key change	Name of the authentication key chain enabled.
Authentication algorithm	Type of authentication algorithm enabled: **hmac** or **md5**
Address families configured	Names of configured address families for the VPN.
Local Address	Address of the local routing device.
Holdtime	Hold time configured with the **hold-time** statement. The hold time is three times the interval at which keepalive messages are sent.
Flags for NLRI inet-label-unicast	Flags related to labeled-unicast: • **TrafficStatistics**—Collection of statistics for labeled-unicast traffic is enabled.
Traffic statistics	Information about labeled-unicast traffic statistics: • **Options**—Options configured for collecting statistics about labeled-unicast traffic. • **File**—Name and location of statistics log files. • **size**—Size of all the log files, in bytes. • **files**—Number of log files.
Traffic Statistics Interval	Time between sample periods for labeled-unicast traffic statistics, in seconds.
Preference	Preference value configured with the **preference** statement.
Number of flaps	Number of times the BGP session has gone down and then come back up.
Peer ID	Router identifier of the peer.

Table 215: show bgp neighbor Output Fields *(continued)*

Field Name	Field Description
Peer Index	Index of this peer in its group.
Local ID	Router identifier of the local routing device.
Local Interface	Name of the interface on the local routing device.
Active holdtime	Hold time the local routing device negotiated with the peer.
Keepalive Interval	Keepalive interval, in seconds.
BFD	Status of BFD failure detection.
Local Address	Name of directly connected interface over which the direct EBGP peering is established.
NLRI for restart configured on peer	Names of address families configured for restart.
NLRI advertised by peer	Address families supported by the peer: **unicast** or **multicast**.
NLRI for this session	Address families being used for this session.
Peer supports Refresh capability	Remote peer's ability to send and request full route table readvertisement (route refresh capability). For more information, see RFC 2918, *Route Refresh Capability for BGP-4*.
Restart time configured on peer	Configured time allowed for restart on the neighbor.
Stale routes from peer are kept for	When graceful restart is negotiated, the maximum time allowed to hold routes from neighbors after the BGP session has gone down.
Restart time requested by this peer	Restart time requested by this neighbor during capability negotiation.
Restart flag received from the peer	When this field appears, the BGP speaker has restarted (Restarting) and this peer should not wait for the **end-of-rib** marker from the speaker before advertising routing information to the speaker.
NLRI that peer supports restart for	Neighbor supports graceful restart for this address family.
NLRI peer can save forwarding state	Neighbor supporting this address family saves all forwarding states.
NLRI that peer saved forwarding for	Neighbor saves all forwarding states for this address family.
NLRI that restart is negotiated for	Router supports graceful restart for this address family.

Table 215: show bgp neighbor Output Fields *(continued)*

Field Name	Field Description
NLRI of received end-of-rib markers	Address families for which end-of-routing-table markers are received from the neighbor.
NLRI of all end-of-rib markers sent	Address families for which end-of-routing-table markers are sent to the neighbor.
Table inet.*number*	Information about the routing table: • **RIB State**—BGP is in the graceful restart process for this routing table: **restart is complete** or **restart in progress**. • **Bit**—Number that represents the entry in the routing table for this peer. • **Send state**—State of the BGP group: **in sync**, **not in sync**, or **not advertising**. • **Active prefixes**—Number of prefixes received from the peer that are active in the routing table. • **Received prefixes**—Total number of prefixes from the peer, both active and inactive, that are in the routing table. • **Accepted prefixes**—Total number of prefixes from the peer that have been accepted by a routing policy. • **Suppressed due to damping**—Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.
Last traffic (seconds)	Last time any traffic was received from the peer or sent to the peer, and the last time the local routing device checked.
Input messages	Messages that BGP has received from the receive socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB.
Output messages	Messages that BGP has written to the transmit socket buffer, showing the total number of messages, number of update messages, number of times a policy is changed and refreshed, and the buffer size in octets. The buffer size is 16 KB.
Output queue	Number of BGP packets that are queued to be transmitted to a particular neighbor for a particular routing table. Output queue **0** is for unicast NLRIs, and queue **1** is for multicast NLRIs.
Trace options	Configured tracing of BGP protocol packets and operations.
Trace file	Name of the file to receive the output of the tracing operation.
Filter Updates recv	(**orf** option only) Number of outbound-route filters received for each configured address family. NOTE: The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list.
Immediate	(**orf** option only) Number of route updates received with the immediate flag set. The immediate flag indicates that the BGP peer should readvertise the updated routes. NOTE: The counter is cumulative. For example, the counter is increased after the remote peer either resends or clears the outbound route filtering prefix list.

Table 215: show bgp neighbor Output Fields *(continued)*

Field Name	Field Description
Filter	(**orf** option only) Type of prefix filter received: **prefix-based** or **extended-community**.
Received filter entries	(**orf** option only) List of received filters displayed.
seq	(**orf** option only) Numerical order assigned to this prefix entry among all the received outbound route filter prefix entries.
prefix	(**orf** option only) Address for the prefix entry that matches the filter.
minlength	(**orf** option only) Minimum prefix length, in bits, required to match this prefix.
maxlength	(**orf** option only) Maximum prefix length, in bits, required to match this prefix.
match	(**orf** option only) For this prefix match, whether to **permit** or **deny** route updates.

show bgp neighbor (CLNS)

```
user@host> show bgp neighbor
Peer: 10.245.245.1+179 AS 200  Local: 10.245.245.3+3770 AS 100
  Type: External    State: Established    Flags: <ImportEval Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Options: <Multihop Preference LocalAddress HoldTime AddressFamily PeerAS
   Rib-group Refresh>
  Address families configured: iso-vpn-unicast
  Local Address: 10.245.245.3 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.245.245.1     Local ID: 10.245.245.3     Active Holdtime: 90
  Keepalive Interval: 30         Peer index: 0
  NLRI advertised by peer: iso-vpn-unicast
  NLRI for this session: iso-vpn-unicast
  Peer supports Refresh capability (2)
  Table bgp.isovpn.0 Bit: 10000
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: in sync
    Active prefixes:              3
    Received prefixes:            3
    Suppressed due to damping:    0
    Advertised prefixes:          3
  Table aaaa.iso.0
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: not advertising
    Active prefixes:              3
    Received prefixes:            3
    Suppressed due to damping:    0
  Last traffic (seconds): Received 6    Sent 5    Checked 5
  Input messages:  Total 1736   Updates 4       Refreshes 0     Octets 33385
  Output messages: Total 1738   Updates 3       Refreshes 0     Octets 33305
  Output Queue[0]: 0
  Output Queue[1]: 0
```

show bgp neighbor (Layer 2 VPN)

```
user@host> show bgp neighbor
Peer: 10.69.103.2     AS 65100 Local: 10.69.103.1     AS 65103
  Type: External    State: Active         Flags: <ImportEval>
  Last State: Idle          Last Event: Start
  Last Error: None
  Export: [ BGP-INET-import ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily   PeerAS
 Refresh>
  Address families configured: inet-unicast
  Local Address: 10.69.103.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
Peer: 10.69.104.2     AS 65100 Local: 10.69.104.1     AS 65104
  Type: External    State: Active         Flags: <ImportEval>
  Last State: Idle          Last Event: Start
  Last Error: None
  Export: [ BGP-L-import ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily PeerAS
 Refresh>
  Address families configured: inet-labeled-unicast
  Local Address: 10.69.104.1 Holdtime: 90 Preference: 170
  Number of flaps: 0
Peer: 10.255.14.182+179 AS 69    Local: 10.255.14.176+2131 AS 69
  Type: Internal    State: Established    Flags: <ImportEval>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
  Rib-group Refresh>
  Address families configured: inet-vpn-unicast l2vpn
  Local Address: 10.255.14.176 Holdtime: 90 Preference: 170
  Number of flaps: 0
  Peer ID: 10.255.14.182    Local ID: 10.255.14.176    Active Holdtime: 90
  Keepalive Interval: 30
  NLRI for restart configured on peer: inet-vpn-unicast l2vpn
  NLRI advertised by peer: inet-vpn-unicast l2vpn
  NLRI for this session: inet-vpn-unicast l2vpn
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 120
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-vpn-unicast l2vpn
  NLRI peer can save forwarding state: inet-vpn-unicast l2vpn
  NLRI that peer saved forwarding for: inet-vpn-unicast l2vpn
  NLRI that restart is negotiated for: inet-vpn-unicast l2vpn
  NLRI of received end-of-rib markers: inet-vpn-unicast l2vpn
  Table bgp.l3vpn.0 Bit: 10000
    RIB State: BGP restart in progress
    RIB State: VPN restart in progress
    Send state: in sync
    Active prefixes:            10
    Received prefixes:          10
    Suppressed due to damping:  0
  Table bgp.l2vpn.0 Bit: 20000
    RIB State: BGP restart in progress
    RIB State: VPN restart in progress
    Send state: in sync
    Active prefixes:            1
    Received prefixes:          1
    Suppressed due to damping:  0
  Table BGP-INET.inet.0 Bit: 30000
    RIB State: BGP restart in progress
    RIB State: VPN restart in progress
```

```
  Send state: in sync
  Active prefixes:            2
  Received prefixes:          2
  Suppressed due to damping:  0
Table BGP-L.inet.0 Bit: 40000
  RIB State: BGP restart in progress
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:            2
  Received prefixes:          2
  Suppressed due to damping:  0
Table LDP.inet.0 Bit: 50000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:            1
  Received prefixes:          1
  Suppressed due to damping:  0
Table OSPF.inet.0 Bit: 60000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:            2
  Received prefixes:          2
  Suppressed due to damping:  0
Table RIP.inet.0 Bit: 70000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:            2
  Received prefixes:          2
  Suppressed due to damping:  0
Table STATIC.inet.0 Bit: 80000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:            1
  Received prefixes:          1
  Suppressed due to damping:  0
Table L2VPN.l2vpn.0 Bit: 90000
  RIB State: BGP restart is complete
  RIB State: VPN restart in progress
  Send state: in sync
  Active prefixes:            1
  Received prefixes:          1
  Suppressed due to damping:  0
Last traffic (seconds): Received 0    Sent 0    Checked 0
Input messages:  Total 14     Updates 13      Refreshes 0     Octets 1053
Output messages: Total 3      Updates 0       Refreshes 0     Octets 105
Output Queue[0]: 0
Output Queue[1]: 0
Output Queue[2]: 0
Output Queue[3]: 0
Output Queue[4]: 0
Output Queue[5]: 0
Output Queue[6]: 0
Output Queue[7]: 0
Output Queue[8]: 0
```

show bgp neighbor (Layer 3 VPN)

```
user@host> show bgp neighbor
```

```
Peer: 4.4.4.4+179     AS 10045 Local: 5.5.5.5+1214    AS 10045
  Type: Internal    State: Established    Flags: <ImportEval>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Export: [ match-all ] Import: [ match-all ]
  Options: <Preference LocalAddress HoldTime GracefulRestart AddressFamily
   Rib-group Refresh>
  Address families configured: inet-vpn-unicast
  Local Address: 5.5.5.5 Holdtime: 90 Preference: 170
  Flags for NLRI inet-labeled-unicast: TrafficStatistics
  Traffic Statistics: Options: all File: /var/log/bstat.log
                                         size 131072 files 10
  Traffic Statistics Interval: 60
  Number of flaps: 0
  Peer ID: 192.168.1.110    Local ID: 192.168.1.111    Active Holdtime: 90
  Keepalive Interval: 30
  NLRI for restart configured on peer: inet-vpn-unicast
  NLRI advertised by peer: inet-vpn-unicast
  NLRI for this session: inet-vpn-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 120
  Stale routes from peer are kept for: 300
  Restart time requested by this peer: 120
  NLRI that peer supports restart for: inet-vpn-unicast
  NLRI peer can save forwarding state: inet-vpn-unicast
  NLRI that peer saved forwarding for: inet-vpn-unicast
  NLRI that restart is negotiated for: inet-vpn-unicast
  NLRI of received end-of-rib markers: inet-vpn-unicast
  NLRI of all end-of-rib markers sent: inet-vpn-unicast
  Table bgp.l3vpn.0 Bit: 10000
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: in sync
    Active prefixes:            2
    Received prefixes:          2
    Suppressed due to damping:  0
  Table vpn-green.inet.0 Bit: 20001
    RIB State: BGP restart is complete
    RIB State: VPN restart is complete
    Send state: in sync
    Active prefixes:            2
    Received prefixes:          2
    Suppressed due to damping:  0
  Last traffic (seconds): Received 15   Sent 20   Checked 20
  Input messages:  Total 40     Updates 2       Refreshes 0     Octets 856
  Output messages: Total 44     Updates 2       Refreshes 0     Octets 1066
  Output Queue[0]: 0
  Output Queue[1]: 0
  Trace options: detail packets
  Trace file: /var/log/bgpgr.log size 131072 files 10
```

show bgp neighbor neighbor-address

```
user@host> show bgp neighbor 192.168.1.111
Peer: 10.255.245.12+179 AS 35  Local: 10.255.245.13+2884 AS 35
  Type: Internal    State: Established  (route reflector client)Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: None
  Options: <Preference LocalAddress HoldTime Cluster AddressFamily Rib-group
  Refresh>
  Address families configured: inet-vpn-unicast inet-labeled-unicast
  Local Address: 10.255.245.13 Holdtime: 90 Preference: 170
  Flags for NLRI inet-vpn-unicast: AggregateLabel
```

```
  Flags for NLRI inet-labeled-unicast: AggregateLabel
  Number of flaps: 0
  Peer ID: 10.255.245.12    Local ID: 10.255.245.13    Active Holdtime: 90
  Keepalive Interval: 30
BFD: disabled
  NLRI advertised by peer: inet-vpn-unicast inet-labeled-unicast
  NLRI for this session: inet-vpn-unicast inet-labeled-unicast
  Peer supports Refresh capability (2)
  Restart time configured on the peer: 300
  Stale routes from peer are kept for: 60
  Restart time requested by this peer: 300
  NLRI that peer supports restart for: inet-unicast inet6-unicast
  NLRI that restart is negotiated for: inet-unicast inet6-unicast
  NLRI of received end-of-rib markers: inet-unicast inet6-unicast
  NLRI of all end-of-rib markers sent: inet-unicast inet6-unicast
  Table inet.0 Bit: 10000
    RIB State: restart is complete
    Send state: in sync
    Active prefixes: 4
    Received prefixes: 6
    Suppressed due to damping: 0
  Table inet6.0 Bit: 20000
    RIB State: restart is complete
    Send state: in sync
    Active prefixes: 0
    Received prefixes: 2
    Suppressed due to damping: 0
  Last traffic (seconds): Received 3    Sent 3    Checked 3
  Input messages:  Total 9      Updates 6       Refreshes 0     Octets 403
  Output messages: Total 7      Updates 3       Refreshes 0     Octets 365
  Output Queue[0]: 0
  Output Queue[1]: 0
  Trace options: detail packets
  Trace file: /var/log/bgpgr size 131072 files 10
```

show bgp neighbor neighbor-address

```
user@host> show bgp neighbor 192.168.4.222
Peer: 192.168.4.222+4902 AS 65501 Local: 192.168.4.221+179 AS 65500
  Type: External    State: Established    Flags: <Sync>
  Last State: OpenConfirm   Last Event: RecvKeepAlive
  Last Error: Cease
  Export: [ export-policy ] Import: [ import-policy ]
  Options: <Preference HoldTime AddressFamily PeerAS PrefixLimit Refresh>
  Address families configured: inet-unicast inet-multicast
  Holdtime: 60000 Preference: 170
  Number of flaps: 4
  Last flap event: RecvUpdate
  Error: 'Cease' Sent: 5 Recv: 0
  Peer ID: 10.255.245.6     Local ID: 10.255.245.5     Active Holdtime: 60000
  Keepalive Interval: 20000      Peer index: 0
  BFD: disabled, down
  Local Interface: fxp0.0
  NLRI advertised by peer: inet-unicast inet-multicast
  NLRI for this session: inet-unicast inet-multicast
  Peer supports Refresh capability (2)
  Table inet.0 Bit: 10000
    RIB State: BGP restart is complete
    Send state: in sync
    Active prefixes:              8
    Received prefixes:            10
    Accepted prefixes:            10
    Suppressed due to damping:    0
```

```
  Advertised prefixes:          3
Table inet.2 Bit: 20000
  RIB State: BGP restart is complete
  Send state: in sync
  Active prefixes:              0
  Received prefixes:            0
  Accepted prefixes:            0
  Suppressed due to damping:    0
  Advertised prefixes:          0
Last traffic (seconds): Received 357  Sent 357  Checked 357
Input messages:  Total 4 Updates 2  Refreshes 0  Octets 211
Output messages: Total 4 Updates 1  Refreshes 0  Octets 147
Output Queue[0]: 0
Output Queue[1]: 0
Trace options:  all
Trace file: /var/log/bgp size 10485760 files 10
```

show bgp neighbor orf neighbor-address detail

```
user@host > show bgp neighbor orf 192.168.165.56 detail
Peer: 192.168.165.56+179 Type: External
  Group: ext1

  inet-unicast
    Filter updates recv:          1 Immediate:          1
    Filter: prefix-based  receive
      Received filter entries:
        seq 1: prefix 2.2.2.2/32: minlen 32: maxlen 32: match deny:

  inet6-unicast
    Filter updates recv:          0 Immediate:          1
    Filter: prefix-based  receive
      Received filter entries:
        *:*
```

show bgp summary

Syntax
show bgp summary
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show bgp summary
<instance *instance-name*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display Border Gateway Protocol (BGP) summary information.

Options
none—Display BGP summary information for all routing instances.

instance *instance-name*—(Optional) Display information for the specified instance only. The instance name can be master for the main instance, or any valid configured instance name or its prefix.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show bgp summary (When a Peer Is Not Established) on page 1826
show bgp summary (When a Peer Is Established) on page 1826
show bgp summary (CLNS) on page 1826
show bgp summary (Layer 2 VPN) on page 1826
show bgp summary (Layer 3 VPN) on page 1827

Output Fields
Table 216 on page 1824 describes the output fields for the **show bgp summary** command. Output fields are listed in the approximate order in which they appear.

Table 216: show bgp summary Output Fields

Field Name	Field Description
Groups	Number of BGP groups.
Peers	Number of BGP peers.
Down peers	Number of down BGP peers.
Table	Name of routing table.
Tot Paths	Total number of paths.
Act Paths	Number of active routes.
Suppressed	Number of routes currently inactive because of damping or other reasons. These routes do not appear in the forwarding table and are not exported by routing protocols.

Table 216: show bgp summary Output Fields *(continued)*

Field Name	Field Description
History	Number of withdrawn routes stored locally to keep track of damping history.
Damp State	Number of routes with a figure of merit greater than zero, but still active because the value has not reached the threshold at which suppression occurs.
Pending	Routes in process by BGP import policy.
Peer	Address of each BGP peer. Each peer has one line of output.
AS	Peer's AS number.
InPkt	Number of packets received from the peer.
OutPkt	Number of packets sent to the peer.
OutQ	Count of the number of BGP packets that are queued to be transmitted to a particular neighbor. It normally is **0** because the queue usually is emptied quickly.
Flaps	Number of times the BGP session has gone down and then come back up.
Last Up/Down	Last time since the neighbor transitioned to or from the established state.
State|#Active /Received/Accepted /Damped	Multipurpose field that displays information about BGP peer sessions. The field's contents depend upon whether a session is established and whether an established session was established in the main routing device or in a routing instance. • If a peer is not established, the field shows the state of the peer session: **Active**, **Connect**, or **Idle**. • If a BGP session is established in the main routing device, the field shows the number of active, received, accepted, and damped routes that are received from a neighbor and appear in the **inet.0** (main) and **inet.2** (multicast) routing tables. For example, **8/10/10/2 and 2/4/4/0** indicate the following: ◦ 8 active routes, 10 received routes, 10 accepted routes, and 2 damped routes from a BGP peer appear in the **inet.0** routing table. ◦ 2 active routes, 4 received routes, 4 accepted routes, and no damped routes from a BGP peer appear in the **inet.2** routing table. • If a BGP session is established in a routing instance, the field indicates the established **(Establ)** state, identifies the specific routing table that receives BGP updates, and shows the number of active, received, and damped routes that are received from a neighbor. For example, **Establ VPN-AB.inet.0: 2/4/0** indicates the following: ◦ The BGP session is established. ◦ Routes are received in the **VPN-AB.inet.0** routing table. ◦ The local routing device has two active routes, four received routes, and no damped routes from a BGP peer. When a BGP session is established, the peers are exchanging update messages.

show bgp summary (When a Peer Is Not Established)

```
user@host> show bgp summary
Groups: 2 Peers: 4 Down peers: 1
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0                 6          4          0          0          0          0
Peer               AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.0.0.3        65002         86         90       0       2       42:54 0/0/0

0/0/0
10.0.0.4        65002         90         91       0       1       42:54 0/2/0

0/0/0
10.0.0.6        65002         87         90       0       3           3 Active
10.1.12.1       65001         89         89       0       1       42:54 4/4/0

0/0/0
```

show bgp summary (When a Peer Is Established)

```
user@host> show bgp summary
Groups: 1 Peers: 3 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
inet.0                 6          4          0          0          0          0
Peer               AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.0.0.2        65002      88675      88652       0       2       42:38 2/4/0

0/0/0
10.0.0.3        65002      54528      54532       0       1     2w4d22h 0/0/0

 0/0/0
10.0.0.4        65002      51597      51584       0       0     2w3d22h 2/2/0

0/0/0
```

show bgp summary (CLNS)

```
user@host> show bgp summary
Groups: 1 Peers: 1 Down peers: 0
Peer               AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.245.245.1      200       1735       1737       0       0    14:26:12 Establ
  bgp.isovpn.0: 3/3/0
  aaaa.iso.0: 3/3/0
```

show bgp summary (Layer 2 VPN)

```
user@host> show bgp summary
Groups: 1 Peers: 5 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State    Pending
bgp.l2vpn.0            1          1          0          0          0          0
inet.0                 0          0          0          0          0          0
Peer               AS      InPkt     OutPkt    OutQ   Flaps Last
Up/Dwn State|#Active/Received/Damped...
10.255.245.35   65299         72         74       0       1       19:00 Establ
  bgp.l2vpn.0: 1/1/0
  frame-vpn.l2vpn.0: 1/1/0
10.255.245.36   65299       2164       2423       0       4       19:50 Establ
  bgp.l2vpn.0: 0/0/0
  frame-vpn.l2vpn.0: 0/0/0
10.255.245.37   65299         36         37       0       4       17:07 Establ
  inet.0: 0/0/0
10.255.245.39   65299        138        168       0       6       53:48 Establ
  bgp.l2vpn.0: 0/0/0
  frame-vpn.l2vpn.0: 0/0/0
```

```
10.255.245.69   65299        134        140       0       6       53:42 Establ
  inet.0: 0/0/0
```

show bgp summary (Layer 3 VPN)

```
user@host> show bgp summary
Groups: 2 Peers: 2 Down peers: 0
Table          Tot Paths  Act Paths Suppressed    History Damp State Pending
bgp.l3vpn.0            2          2          0          0          0       0
Peer                AS      InPkt     OutPkt    OutQ   Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.39.1.5            2         21         22       0       0        6:26 Establ
  VPN-AB.inet.0: 1/1/0
10.255.71.15         1         19         21       0       0        6:17 Establ
  bgp.l3vpn.0: 2/2/0
  VPN-A.inet.0: 1/1/0
  VPN-AB.inet.0: 2/2/0
  VPN-B.inet.0: 1/1/0
```

show ipv6 neighbors

Syntax show ipv6 neighbors

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display information about the IPv6 neighbor cache.

Options This command has no options.

Required Privilege Level view

Related Documentation
- **clear ipv6 neighbors on page 990**

List of Sample Output
show ipv6 neighbors on page 1828
show ipv6 neighbors on page 1828

Output Fields Table 217 on page 1828 describes the output fields for the **show ipv6 neighbors** command. Output fields are listed in the approximate order in which they appear.

Table 217: show ipv6 neighbors Output Fields

Field Name	Field Description
IPv6 Address	Name of the IPv6 interface.
Linklayer Address	Link-layer address.
State	State of the link: **up**, **down**, **incomplete**, **reachable**, **stale**, or **unreachable**.
Exp	Number of seconds until the entry expires.
Rtr	Whether the neighbor is a routing device: **yes** or **no**.
Secure	Whether this entry was created using the Secure Neighbor Discovery (SEND) protocol: **yes** or **no**.
Interface	Name of the interface.

show ipv6 neighbors

```
user@host> show ipv6 neighbors
IPv6 Address                 Linklayer Address  State       Exp  Rtr  Interface
fe80::2a0:c9ff:fe5b:4c1e     00:a0:c9:5b:4c:1e  reachable   15   yes  fxp0.0
```

show ipv6 neighbors

```
user@host > show ipv6 neighbors
IPv6 Address                 Linklayer Address  State       Exp Rtr Secure
Interface
```

```
fe80::14fb:5dcf:54bd:ff76    00:90:69:a0:a8:bc  stale       1113 yes yes
ge-3/2/0.0
```

show isis adjacency

Syntax
show isis adjacency
<brief | detail | extensive>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show isis adjacency
<brief | detail | extensive>
<instance *instance-name*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display information about Intermediate System-to-Intermediate System (IS-IS) neighbors.

Options
none—Display standard information about IS-IS neighbors for all routing instances.

brief | detail | extensive—(Optional) Display the specified level of output.

instance *instance-name*—(Optional) Display adjacencies for the specified routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- **clear isis adjacency on page 1760**

List of Sample Output
show isis adjacency on page 1832
show isis adjacency brief on page 1832
show isis adjacency detail on page 1832
show isis adjacency extensive on page 1833

Output Fields
Table 218 on page 1830 describes the output fields for the **show isis adjacency** command. Output fields are listed in the approximate order in which they appear.

Table 218: show isis adjacency Output Fields

Field Name	Field Description	Level of Output
Interface	Interface through which the neighbor is reachable.	All levels
System	System identifier (**sysid**), displayed as a name, if possible.	**brief**

Table 218: show isis adjacency Output Fields *(continued)*

Field Name	Field Description	Level of Output
L or Level	Level: • **1**—Level 1 only • **2**—Level 2 only • **3**—Level 1 and Level 2 An exclamation point (**!**) preceding the level number indicates that the adjacency is missing an IP address.	All levels
State	State of the adjacency: **Up**, **Down**, **New**, **One-way**, **Initializing**, or **Rejected**.	All levels
Hold (secs)	Remaining hold time of the adjacency.	**brief**
SNPA	Subnetwork point of attachment (MAC address of the next hop).	**brief**
Expires in	How long until the adjacency expires, in seconds.	**detail**
Priority	Priority to become the designated intermediate system.	**detail extensive**
Up/Down transitions	Count of adjacency status changes from **Up** to **Down** or from **Down** to **Up**.	**detail**
Last transition	Time of the last **Up/Down** transition.	**detail**
Circuit type	Bit mask of levels on this interface: **L1**=Level 1 router; **L2**=Level 2 router; **L1/L2**=both Level 1 and Level 2 router.	**detail**
Speaks	Protocols supported by this neighbor.	**detail extensive**
MAC address	MAC address of the interface.	**detail extensive**
Topologies	Supported topologies.	**detail extensive**
Restart capable	Whether a neighbor is capable of graceful restart: **Yes** or **No**.	**detail extensive**
Adjacency advertisement: Advertise	This router has signaled not to advertise this interface to its neighbors in their label-switched paths (LSPs).	**detail extensive**
Adjacency advertisement: Suppress	This neighbor has signaled not to advertise the interface in the router's outbound LSPs.	**detail extensive**
IP addresses	IP address of this neighbor.	**detail extensive**

Table 218: show isis adjacency Output Fields *(continued)*

Field Name	Field Description	Level of Output

Transition log	List of recent transitions, including:
	• **When**—Time at which an IS-IS adjacency transition occurred.
	• **State**—Current state of the IS-IS adjacency (**up**, **down**, or **rejected**).
	  • **Up**—Adjacency is up and operational.
	  • **Down**—Adjacency is down and not available.
	  • **Rejected**—Adjacency has been rejected.
	• **Event**—Type of transition that occurred.
	  • **Seenself**—Possible routing loop has been detected.
	  • **Interface down**—IS-IS interface has gone down and is no longer available.
	  • **Error**—Adjacency error.
	• **Down reason**—Reason that an IS-IS adjacency is down:
	  • **3-Way Handshake Failed**—Connection establishment failed.
	  • **Address Mismatch**—Address mismatch caused link failure.
	  • **Aged Out**—Link expired.
	  • **ISO Area Mismatch**—IS-IS area mismatch caused link failure.
	  • **Bad Hello**—Unacceptable hello message caused link failure.
	  • **BFD Session Down**—Bidirectional failure detection caused link failure.
	  • **Interface Disabled**—IS-IS interface is disabled.
	  • **Interface Down**—IS-IS interface is unavailable.
	  • **Interface Level Disabled**—IS-IS level is disabled.
	  • **Level Changed**—IS-IS level has changed on the adjacency.
	  • **Level Mismatch**—Levels on adjacency are not compatible.
	  • **MPLS LSP Down**—Label-switched path (LSP) is unavailable.
	  • **MT Topology Changed**—IS-IS topology has changed.
	  • **MT Topology Mismatch**—IS-IS topology is mismatched.
	  • **Remote System ID Changed**—Adjacency peer system ID changed.
	  • **Protocol Shutdown**—IS-IS protocol is disabled.
	  • **CLI Command**—Adjacency brought down by user.
	  • **Unknown**—Unknown.
		extensive

show isis adjacency

```
user@host> show isis adjacency
Interface             System         L State       Hold (secs) SNPA
at-2/3/0.0            ranier         3  Up                  23
```

show isis adjacency brief

The output for the **show isis adjacency brief** command is identical to that for the **show isis adjacency** command. For sample output, see **show isis adjacency on page 1832**.

show isis adjacency detail

```
user@host> show isis adjacency detail
ranier
  Interface: at-2/3/0.0, Level: 3, State: Up, Expires in 21 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 00:01:09 ago
  Circuit type: 3, Speaks: IP, IPv6
  Topologies: Unicast
```

```
  Restart capable: Yes
  IP addresses: 11.1.1.2
```

show isis adjacency extensive

```
user@host> show isis adjacency extensive
ranier
  Interface: at-2/3/0.0, Level: 3, State: Up, Expires in 22 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 00:01:16 ago
  Circuit type: 3, Speaks: IP, IPv6
  Topologies: Unicast
  Restart capable: Yes
  IP addresses: 11.1.1.2
  Transition log:
  When                  State        Event           Down reason
  Wed Nov  8 21:24:25   Up           Seenself
```

show isis authentication

Syntax

show isis authentication
<brief | detail | extensive>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)

show isis authentication
<brief | detail | extensive>
<instance *instance-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display information about Intermediate System-to-Intermediate System (IS-IS) authentication.

Options

none—Display information about IS-IS authentication.

brief | detail | extensive—(Optional) Display the specified level of output.

instance *instance-name*—(Optional) Display IS-IS authentication for the specified routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show isis authentication on page 1835

Output Fields

Table 219 on page 1834 describes the output fields for the **show isis authentication** command. Output fields are listed in the approximate order in which they appear.

Table 219: show isis authentication Output Fields

Field Name	Field Description
Interface	Interface name.
Level	IS-IS level.
IIH Auth	IS-IS Hello (IIH) packet authentication type.
CSN Auth	Complete sequence number authentication type.
PSN Auth	Partial sequence number authentication type.
L1 LSP Authentication	Layer 1 link-state PDU authentication type.

Table 219: show isis authentication Output Fields *(continued)*

Field Name	Field Description
L2 LSP Authentication	Layer 2 link-state PDU authentication type.

show isis authentication

```
user@host> show isis authentication
Interface             Level IIH Auth  CSN Auth  PSN Auth
at-2/3/0.0            1     Simple    Simple    Simple
                      2     MD5       MD5       MD5

L1 LSP Authentication: Simple
L2 LSP Authentication: MD5
```

show isis backup coverage

Syntax

show isis backup coverage
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)

show isis backup coverage
<instance *instance-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display information about the level of backup coverage available.

Options

none—Display information about the level of backup coverage available for all the nodes and prefixes in the network.

instance *instance-name*—(Optional) Display information about the level of backup coverage for a specific IS-IS routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- **show isis backup label-switched-path on page 1838**

List of Sample Output

show isis backup coverage on page 1837

Output Fields

Table 220 on page 1836 lists the output fields for the **show isis backup coverage** command. Output fields are listed in the approximate order in which they appear.

Table 220: show isis backup coverage Output Fields

Field Name	Field Description
Topology	Type of topology or address family: **IPV4 Unicast** or **IPV6 Unicast**.
Level	IS-IS level: • **1**—Level 1 • **2**—Level 2
Node	By topology, the percentage of all routes configured on the node that are protected through backup coverage.
IPv4 Unicast	Percentage of IPv4 unicast routes that are protected through backup coverage.
IPv6 Unicast	Percentage of IPv6 unicast routes that are protected through backup coverage.

Table 220: show isis backup coverage Output Fields *(continued)*

Field Name	Field Description
CLNS	Percentage of Connectionless Network Service (CLNS) routes that are protected through backup coverage.

show isis backup coverage

```
user@host> show isis backup coverage
Backup Coverage:
    Topology         Level   Node    IPv4    IPv6    CLNS
    IPV4 Unicast         2  28.57%  22.22%   0.00%   0.00%
    IPV6 Unicast         2   0.00%   0.00%   0.00%   0.00%
```

show isis backup label-switched-path

Syntax **show isis backup label-switched-path**
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch) **show isis backup label-switched-path**

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display information about MPLS label-switched-paths (LSPs) designated as backup routes for IS-IS routes.

Options none—Display information about MPLS LSPs designated as backup routes for IS-IS routes.

logical-system (all | *logical-system-name*—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level view

Related Documentation
- **show isis backup coverage on page 1836**

List of Sample Output **show isis backup label-switched-path on page 1839**

Output Fields Table 221 on page 1838 lists the output fields for the **show isis backup label-switched-path** command. Output fields are listed in the approximate order in which they appear.

Table 221: show isis backup label-switched-path Output Fields

Field Name	Field Description
Backup MPLS LSPs	List of MPLS LSPs designated as backup paths for IS-IS routes.
Egress	IP address of the egress routing device for the LSP.

Status	State of the LSP:
	• **Up**—The router can detect RSVP hello messages from the neighbor.
	• **Down**—The router has received one of the following indications:
	  ◦ Communication failure from the neighbor.
	  ◦ Communication from IGP that the neighbor is unavailable.
	  ◦ Change in the sequence numbers in the RSVP hello messages sent by the neighbor.
	• **Deleted**—LSP is no longer available as a backup path.

Last change	Time elapsed since the neighbor state changed either from up or down or from down to up. The format is ***hh*:*mm*:*ss***.
TE-metric	Configured traffic engineering metric.

Table 221: show isis backup label-switched-path Output Fields *(continued)*

Field Name	Field Description
Metric	Configured metric.

show isis backup label-switched-path

```
user@host> show isis backup label-switched-path
 Backup MPLS LSPs:
 f-to-g, Egress: 192.168.1.4, Status: up, Last change: 06:12:03
    TE-metric: 9, Metric: 0
```

show isis backup spf results

Syntax

show isis backup spf results
<instance *instance-name*>
<level (1 | 2)>
<logical-system (all | *logical-system-name*)>
<no-coverage>
<topology (ipv4-unicast | ipv6-multicast | ipv6-unicast | unicast)>

Syntax (J-EX Series Switch)

show isis backup spf results
<instance *instance-name*>
<level (1 | 2)>
<no-coverage>
<topology (ipv4-unicast | unicast)>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display information about IS-IS shortest-path-first (SPF) calculations for backup paths.

Options

none—Display information about IS-IS shortest-path-first (SPF) calculations for all backup paths for all destination nodes.

instance *instance-name*—(Optional) Display SPF calculations for backup paths for the specified routing instance.

level (1 | 2)—(Optional) Display SPF calculations for the backup paths for the specified IS-IS level.

logical-system *logical-system-name*—(Optional) Display SPF calculations for the backup paths for all logical systems or on a particular logical system.

no-coverage—(Optional) Display SPF calculations only for destinations that do not have backup coverage.

topology (ipv4-multicast | ipv6-multicast | ipv6-unicast | unicast)—(Optional) Display SPF calculations for backup paths for the specified topology only.

Required Privilege Level

view

Related Documentation

- **show isis backup coverage on page 1836**

List of Sample Output

show isis backup spf results on page 1841

Output Fields

Table 222 on page 1840 lists the output fields for the **show isis backup spf results** command. Output fields are listed in the approximate order in which they appear.

Table 222: show isis backup spf results Output Fields

Field Name	Field Description
node-name	Name of the destination node.

Table 222: show isis backup spf results Output Fields *(continued)*

Field Name	Field Description
Address	Address of the destination node.
Primary next-hop	Interface and name of the node of the primary next hop to reach the destination.
Root	Name of the next-hop neighbor.
Metric	Metric to the node.
Eligible	Indicates that the next-hop neighbor has been designated as a backup path to the destination node.
Backup next-hop	Name of the interface of the backup next hop.
SNPA	Subnetwork point of attachment (MAC address of the next hop).
LSP	Name of the MPLS LSP designated as a backup path.
Not eligible	Indicates that the next-hop neighbor cannot function as a backup path to the destination.
Reason	Describes why the next-hop neighbor is designated as **Not eligible** as a backup path.

show isis backup spf results

```
user@host> show isis backup spf results
IS-IS level 1 SPF results:
  0 nodes

 IS-IS level 2 SPF results:
kobuk.00, Address 0x8d85600
  Primary next-hop: ge-0/2/0.0, camaro, SNPA:  0:90:69:f:62:fa
  Primary next-hop: so-0/1/2.0, crater
  Primary next-hop: ge-0/2/0.0, camaro, SNPA:  0:90:69:f:62:fa
  Primary next-hop: so-0/1/2.0, crater
    Root: crater, Metric: 10
      Not eligible, Reason: Primary next-hop multipath
    Root: camaro, Metric: 10
      Not eligible, Reason: Primary next-hop multipath
    Root: olympic, Metric: 25
      Not eligible, Reason: Primary next-hop multipath
glacier.00, Address 0x8d85200
  Primary next-hop: so-0/1/2.0, crater
  Primary next-hop: so-0/1/2.0, crater
    Root: crater, Metric: 10
      Not eligible, Reason: Primary next-hop link fate sharing
    Root: olympic, Metric: 15
      Eligible, Backup next-hop: ge-0/2/0.0, camaro, SNPA:  0:90:69:f:62:fa
      Eligible, Backup next-hop: so-1/0/2.0, olympic
      Eligible, Backup next-hop: ge-0/2/0.0, camaro, SNPA:  0:90:69:f:62:fa
      Eligible, Backup next-hop: so-1/0/2.0, olympic
    Root: camaro, Metric: 20
```

```
      Eligible, Backup next-hop: ge-0/2/0.0, camaro, SNPA:  0:90:69:f:62:fa
      Eligible, Backup next-hop: so-1/0/2.0, olympic
      Eligible, Backup next-hop: ge-0/2/0.0, camaro, SNPA:  0:90:69:f:62:fa
      Eligible, Backup next-hop: so-1/0/2.0, olympic
olympic.00, Address 0x8d00c00
  Primary next-hop: so-1/0/2.0, olympic
  Primary next-hop: so-1/0/2.0, olympic
    Root: olympic, Metric: 0
      Not eligible, Reason: Primary next-hop link fate sharing
    Root: crater, Metric: 20
      track-item: olympic.00-00
      track-item: banff.00-00
      Not eligible, Reason: Path loops
    Root: camaro, Metric: 20
      track-item: olympic.00-00
      track-item: banff.00-00
      Not eligible, Reason: Path loops
camaro.00, Address 0x8d85a00
  Primary next-hop: ge-0/2/0.0, camaro, SNPA:  0:90:69:f:62:fa
  Primary next-hop: ge-0/2/0.0, camaro, SNPA:  0:90:69:f:62:fa
    Root: camaro, Metric: 0
      Not eligible, Reason: Primary next-hop link fate sharing
    Root: crater, Metric: 20
      track-item: camaro.00-00
      track-item: banff.00-00
      Not eligible, Reason: Path loops
    Root: olympic, Metric: 20
      track-item: camaro.00-00
      track-item: banff.00-00
      Not eligible, Reason: Path loops
crater.00, Address 0x8d85000
  Primary next-hop: so-0/1/2.0, crater
  Primary next-hop: so-0/1/2.0, crater
    Root: crater, Metric: 0
      Not eligible, Reason: Primary next-hop link fate sharing
    Root: camaro, Metric: 20
      track-item: crater.00-00
      track-item: banff.00-00
      Not eligible, Reason: Path loops
    Root: olympic, Metric: 20
      track-item: crater.00-00
      track-item: banff.00-00
      Not eligible, Reason: Path loops
  5 nodes
```

show isis database

Syntax

show isis database
<brief | detail | extensive>
<instance *instance-name*>
<level (1 | 2)>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)

show isis database
<brief | detail | extensive>
<level (1 | 2)>
<instance *instance-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display the entries in the Intermediate System-to-Intermediate System (IS-IS) link-state database, which contains data about PDU packets.

Options

none—Display standard information about IS-IS link-state database entries for all routing instances.

brief | detail | extensive—(Optional) Display the specified level of output.

instance *instance-name*—(Optional) Display entries for the specified routing instance.

level (1 | 2)—(Optional) Display entries for the specified IS-IS level.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- **clear isis database on page 1762**

List of Sample Output

show isis database on page 1845
show isis database brief on page 1845
show isis database detail on page 1846
show isis database extensive on page 1847
show isis database extensive (CLNS) on page 1848

Output Fields

Table 223 on page 1844 describes the output fields for the **show isis database** command. Output fields are listed in the approximate order in which they appear. Fields that contain internal IS-IS information useful only in troubleshooting obscure problems are not described in the table. For more details about these fields, contact your customer support representative.

Table 223: show isis database Output Fields

Field Name	Field Description	Level of Output
Interface name	Name of the interface on which the LSP has been received; always **IS-IS** for this command.	All levels
level	Level of intermediate system: • **1**—Intermediate system routes within an area; when the destination is outside an area, it routes toward a Level 2 system. • **2**—Intermediate system routes between areas and toward other ASs.	All levels
LSP ID	Link-state PDU identifier.	All levels
Sequence	Sequence number of the link-state PDU.	All levels
Checksum	Checksum value of the link-state PDU.	All levels
Lifetime (secs)	Remaining lifetime of the link-state PDU, in seconds.	All levels
Attributes	Attributes of the specified database: **L1**, **L2**, **Overload**, or **Attached** (L1 only).	none **brief**
# LSPs	Total number of LSPs in the specified link-state database.	none **brief**
IP prefix	Prefix advertised by this link-state PDU.	**detail extensive**
IS neighbor	IS-IS neighbor of the advertising system.	**detail extensive**
ES neighbor	(J Series routers only) An ES-IS neighbor of the advertising system.	**detail extensive**
IP prefix	IPv4 prefix advertised by this link-state PDU.	**detail extensive**
V6 prefix	IPv6 prefix advertised by this link-state PDU.	**detail extensive**
Metric	Metric of the prefix or neighbor.	**detail extensive**
Header	• **LSP ID**—Link state PDU identifier of the header. • **Length**—Header length. • **Allocated Length**—Amount of length available for the header. • **Router ID**—Address of the local routing device. • **Remaining Lifetime**—Remaining lifetime of the link-state PDU, in seconds.	**extensive**

Table 223: show isis database Output Fields *(continued)*

Field Name	Field Description	Level of Output
Packet	• **LSP ID**—The identifier for the link-state packet. • **Length**—Packet length. • **Lifetime**—Remaining lifetime, in seconds. • **Checksum**—The checksum of the LSP. • **Sequence**—The sequence number of the LSP. Every time the LSP is updated, this number increments. • **Attributes**—Packet attributes. • **NLPID**—Network layer protocol identifier. • **Fixed length**—Specifies the set length for the packet.	**extensive**
TLVs	• **Area Address**—Area addresses that the routing device can reach. • **Speaks**—Supported routing protocols. • **IP router id**—ID of the routing device (usually the IP address). • **IP address**—IPv4 address. • **Hostname**—Assigned name of the routing device. • **IP prefix**—IP prefix of the routing device. • **Metric**—IS-IS metric that measures the cost of the adjacency between the originating routing device and the advertised routing device. • **IP extended prefix**—Extended IP prefix of the routing device. • **IS neighbor**—Directly attached neighbor's name and metric. • **IS extended neighbor**—Directly attached neighbor's name, metric, and IP address.	**extensive**

show isis database

```
user@host> show isis database
IS-IS level 1 link-state database:
LSP ID                  Sequence Checksum Lifetime Attributes
kobuk.00-00                  0x3   0x3167     1057 L1 L2
camaro.00-00                 0x5   0x770e     1091 L1 L2
ranier.00-00                 0x4   0xaa95     1091 L1 L2
glacier.00-00                0x4   0x206f     1089 L1 L2
glacier.02-00                0x1   0xd141     1089 L1 L2
badlands.00-00               0x3   0x87a2     1093 L1 L2
  6 LSPs

IS-IS level 2 link-state database:
LSP ID                  Sequence Checksum Lifetime Attributes
kobuk.00-00                  0x6   0x8d6b     1096 L1 L2
camaro.00-00                 0x9   0x877b     1101 L1 L2
ranier.00-00                 0x8   0x855d     1103 L1 L2
glacier.00-00                0x7   0xf892     1098 L1 L2
glacier.02-00                0x1   0xd141     1089 L1 L2
badlands.00-00               0x6    0x562     1105 L1 L2
  6 LSPs
```

show isis database brief

The output for the **show isis database brief** command is identical to that for the **show isis database** command. For sample output, see **show isis database on page 1845**.

show isis database detail

```
user@host> show isis database detail
IS-IS level 1 link-state database:

kobuk.00-00 Sequence: 0x3, Checksum: 0x3167, Lifetime: 1048 secs
   IS neighbor: glacier.00                    Metric:       10
   IP prefix: 10.255.70.103/32                Metric:        0 Internal Up
   IP prefix: 43.1.1.0/24                     Metric:       10 Internal Up
   V6 prefix: abcd::10:255:70:103/128         Metric:        0 Internal Up

camaro.00-00 Sequence: 0x5, Checksum: 0x770e, Lifetime: 1082 secs
   IS neighbor: ranier.00                     Metric:       10
   IS neighbor: glacier.02                    Metric:       10
   IP prefix: 10.255.71.52/32                 Metric:        0 Internal Up
   IP prefix: 23.1.1.0/24                     Metric:       10 Internal Up
   IP prefix: 34.1.1.0/24                     Metric:       10 Internal Up
   V6 prefix: abcd::10:255:71:52/128          Metric:        0 Internal Up

ranier.00-00 Sequence: 0x4, Checksum: 0xaa95, Lifetime: 1082 secs
   IS neighbor: camaro.00                     Metric:       10
   IS neighbor: badlands.00                   Metric:       10
   IP prefix: 10.255.71.241/32                Metric:        0 Internal Up
   IP prefix: 11.1.1.0/24                     Metric:       10 Internal Up
   IP prefix: 23.1.1.0/24                     Metric:       10 Internal Up
   V6 prefix: abcd::10:255:71:241/128         Metric:        0 Internal Up

glacier.00-00 Sequence: 0x4, Checksum: 0x206f, Lifetime: 1080 secs
   IS neighbor: kobuk.00                      Metric:       10
   IS neighbor: glacier.02                    Metric:       10
   IP prefix: 10.255.71.242/32                Metric:        0 Internal Up
   IP prefix: 34.1.1.0/24                     Metric:       10 Internal Up
   IP prefix: 43.1.1.0/24                     Metric:       10 Internal Up
   V6 prefix: abcd::10:255:71:242/128         Metric:        0 Internal Up

glacier.02-00 Sequence: 0x1, Checksum: 0xd141, Lifetime: 1080 secs
   IS neighbor: camaro.00                     Metric:        0
   IS neighbor: glacier.00                    Metric:        0

badlands.00-00 Sequence: 0x3, Checksum: 0x87a2, Lifetime: 1084 secs
   IS neighbor: ranier.00                     Metric:       10
   IP prefix: 10.255.71.244/32                Metric:        0 Internal Up
   IP prefix: 11.1.1.0/24                     Metric:       10 Internal Up
   V6 prefix: abcd::10:255:71:244/128         Metric:        0 Internal Up

IS-IS level 2 link-state database:

kobuk.00-00 Sequence: 0x6, Checksum: 0x8d6b, Lifetime: 1088 secs
   IS neighbor: glacier.00                    Metric:       10
   IP prefix: 10.255.70.103/32                Metric:        0 Internal Up
   IP prefix: 10.255.71.52/32                 Metric:       20 Internal Up
   IP prefix: 10.255.71.241/32                Metric:       30 Internal Up
   IP prefix: 10.255.71.242/32                Metric:       10 Internal Up
   IP prefix: 10.255.71.244/32                Metric:       40 Internal Up
   IP prefix: 11.1.1.0/24                     Metric:       40 Internal Up
   IP prefix: 23.1.1.0/24                     Metric:       30 Internal Up
   IP prefix: 34.1.1.0/24                     Metric:       20 Internal Up
   IP prefix: 43.1.1.0/24                     Metric:       10 Internal Up
   V6 prefix: abcd::10:255:70:103/128         Metric:        0 Internal Up

camaro.00-00 Sequence: 0x9, Checksum: 0x877b, Lifetime: 1092 secs
   IS neighbor: ranier.00                     Metric:       10
   IS neighbor: glacier.02                    Metric:       10
```

```
   IP prefix: 10.255.70.103/32               Metric:       20 Internal Up
   IP prefix: 10.255.71.52/32                Metric:        0 Internal Up
   IP prefix: 10.255.71.241/32               Metric:       10 Internal Up
   IP prefix: 10.255.71.242/32               Metric:       10 Internal Up
   IP prefix: 10.255.71.244/32               Metric:       20 Internal Up
   IP prefix: 11.1.1.0/24                    Metric:       20 Internal Up
   IP prefix: 23.1.1.0/24                    Metric:       10 Internal Up
   IP prefix: 34.1.1.0/24                    Metric:       10 Internal Up
   IP prefix: 43.1.1.0/24                    Metric:       20 Internal Up
   V6 prefix: abcd::10:255:71:52/128         Metric:        0 Internal Up

ranier.00-00 Sequence: 0x8, Checksum: 0x855d, Lifetime: 1094 secs
   IS neighbor: camaro.00                    Metric:       10
   IS neighbor: badlands.00                  Metric:       10
   IP prefix: 10.255.70.103/32               Metric:       30 Internal Up
   IP prefix: 10.255.71.52/32                Metric:       10 Internal Up
   IP prefix: 10.255.71.241/32               Metric:        0 Internal Up
   IP prefix: 10.255.71.242/32               Metric:       20 Internal Up
   IP prefix: 10.255.71.244/32               Metric:       10 Internal Up
   IP prefix: 11.1.1.0/24                    Metric:       10 Internal Up
   IP prefix: 23.1.1.0/24                    Metric:       10 Internal Up
   IP prefix: 34.1.1.0/24                    Metric:       20 Internal Up
   IP prefix: 43.1.1.0/24                    Metric:       30 Internal Up
   V6 prefix: abcd::10:255:71:241/128        Metric:        0 Internal Up

glacier.00-00 Sequence: 0x7, Checksum: 0xf892, Lifetime: 1089 secs
   IS neighbor: kobuk.00                     Metric:       10
   IS neighbor: glacier.02                   Metric:       10
   IP prefix: 10.255.70.103/32               Metric:       10 Internal Up
   IP prefix: 10.255.71.52/32                Metric:       10 Internal Up
   IP prefix: 10.255.71.241/32               Metric:       20 Internal Up
   IP prefix: 10.255.71.242/32               Metric:        0 Internal Up
   IP prefix: 10.255.71.244/32               Metric:       30 Internal Up
   IP prefix: 11.1.1.0/24                    Metric:       30 Internal Up
   IP prefix: 23.1.1.0/24                    Metric:       20 Internal Up
   IP prefix: 34.1.1.0/24                    Metric:       10 Internal Up
   IP prefix: 43.1.1.0/24                    Metric:       10 Internal Up
   V6 prefix: abcd::10:255:71:242/128        Metric:        0 Internal Up

glacier.02-00 Sequence: 0x1, Checksum: 0xd141, Lifetime: 1080 secs
   IS neighbor: camaro.00                    Metric:        0
   IS neighbor: glacier.00                   Metric:        0

badlands.00-00 Sequence: 0x6, Checksum: 0x562, Lifetime: 1096 secs
   IS neighbor: ranier.00                    Metric:       10
   IP prefix: 10.255.70.103/32               Metric:       40 Internal Up
   IP prefix: 10.255.71.52/32                Metric:       20 Internal Up
   IP prefix: 10.255.71.241/32               Metric:       10 Internal Up
   IP prefix: 10.255.71.242/32               Metric:       30 Internal Up
   IP prefix: 10.255.71.244/32               Metric:        0 Internal Up
   IP prefix: 11.1.1.0/24                    Metric:       10 Internal Up
   IP prefix: 23.1.1.0/24                    Metric:       20 Internal Up
   IP prefix: 34.1.1.0/24                    Metric:       30 Internal Up
   IP prefix: 43.1.1.0/24                    Metric:       40 Internal Up
   V6 prefix: abcd::10:255:71:244/128        Metric:        0 Internal Up
```

show isis database extensive

```
user@host> show isis database extensive isis2
IS-IS level 1 link-state database:

IS-IS level 2 link-state database:
```

```
isis2.00-00  Sequence: 0x82, Checksum: 0x6cc3, Lifetime: 1126 secs
   IS neighbor:                      isis1.00  Metric:      10
   IS neighbor:                      isis3.00  Metric:      10
   IP prefix:               10.255.245.202/32 Metric:       0 Internal
   IP prefix:                 192.168.36.0/29 Metric:      10 Internal
   IP prefix:                192.168.36.16/30 Metric:      10 Internal
   IP prefix:                192.168.36.24/30 Metric:      10 Internal

  Header: LSP ID: isis2.00-00, Length: 234 bytes
    Allocated length: 234 bytes, Router ID: 10.255.245.202
    Remaining lifetime: 1126 secs, Level: 2,Interface: 4
    Estimated free bytes: 0, Actual free bytes: 0
    Aging timer expires in: 1126 secs
    Protocols: IP, IPv6

  Packet: LSP ID: isis2.00-00, Length: 234 bytes, Lifetime : 1198 secs
    Checksum: 0x6cc3, Sequence: 0x82, Attributes: 0x3 <L1 L2>
    NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
    Packet type: 20, Packet version: 1, Max area: 0

  TLVs:
    Area address: 47.0005.80ff.f800.0000.0108.0001 (13)
    Speaks: IP
    Speaks: IPv6
    IP router id: 10.255.245.202
    IP address: 10.255.245.202
    Hostname: isis2
    IS neighbor: isis3.00, Internal, Metric: default 10
    IS neighbor: isis1.00, Internal, Metric: default 10
    IS neighbor: isis3.00, Metric: default 10
      IP address: 192.168.36.25
      Neighbor's IP address: 192.168.36.26
    IS neighbor: isis1.00, Metric: default 10
      IP address: 192.168.36.18
      Neighbor's IP address: 192.168.36.17
    IP prefix: 10.255.245.202/32, Internal, Metric: default 0
    IP prefix: 192.168.36.0/29, Internal, Metric: default 10
    IP prefix: 192.168.36.24/30, Internal, Metric: default 10
    IP prefix: 192.168.36.16/30, Internal, Metric: default 10
    IP prefix: 10.255.245.202/32 metric 0 up
      6 bytes of subtlvs
      Administrative tag 1: 1000
    IP prefix: 192.168.36.0/29 metric 10 up
    IP prefix: 192.168.36.24/30 metric 10 up
    IP prefix: 192.168.36.16/30 metric 10 up
  No queued transmissions
```

show isis database extensive (CLNS)

```
user@host> show isis database extensive
IS-IS level 1 link-state database:
isis2.00-00 Sequence: 0x1256, Checksum: 0x53da, Lifetime: 582 secs
   IS neighbor: pro1-a.02                    Metric:       10
   ES neighbor: toothache                    Metric:        0
   ES neighbor: 1921.6800.4002               Metric:       10
   IP prefix: 192.168.37.64/29               Metric:       10 Internal Up

  Header: LSP ID: toothache.00-00, Length: 140 bytes
    Allocated length: 284 bytes, Router ID: 0.0.0.0
    Remaining lifetime: 582 secs, Level: 1,Interface: 66
    Estimated free bytes: 144, Actual free bytes: 144
    Aging timer expires in: 582 secs
    Protocols: IP, CLNS
```

```
Packet: LSP ID: toothache.00-00, Length: 140 bytes, Lifetime : 1199 secs
  Checksum: 0x53da, Sequence: 0x1256, Attributes: 0xb <L1 L2 Attached>
  NLPID: 0x83, Fixed length: 27 bytes, Version: 1, Sysid length: 0 bytes
  Packet type: 18, Packet version: 1, Max area: 0

TLVs:
  Area address: 47.0005.80ff.f800.0000.0108.0001 (13)
  Speaks: CLNP
  Speaks: IP
  Hostname: toothache
  IP address: 192.168.37.69
  IP extended prefix: 192.168.37.64/29 metric 10 up
  IP prefix: 192.168.37.64/29, Internal, Metric: default 10, Up
  IS neighbor: pro1-a.02, Internal, Metric: default 10
  IS extended neighbor: pro1-a.02, Metric: default 10
  ES neighbor TLV: Internal, Metric: default 0
    ES: toothache
  ES neighbor TLV: Internal, Metric: default 10
    ES: 1921.6800.4002
No queued transmissions
```

show isis hostname

Syntax
show isis hostname
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show isis hostname

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display Intermediate System-to-Intermediate System (IS-IS) hostname database information.

Options
none—Display IS-IS hostname database information.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show isis hostname on page 1850

Output Fields
Table 224 on page 1850 describes the output fields for the **show isis hostname** command. Output fields are listed in the approximate order in which they appear.

Table 224: show isis hostname Output Fields

Field Name	Field Description
System Id	System identifier mapped to the hostname.
Hostname	Hostname mapped to the system identifier.
Type	Type of mapping between system identifier and hostname. • **Dynamic**—Hostname mapping determined as described in RFC 2763, *Dynamic Hostname Exchange Mechanism for IS-IS*. • **Static**—Hostname mapping configured by user.

show isis hostname

```
user@host> show isis hostname
   IS-IS hostname database:
   System Id      Hostname                                        Type
   1921.6800.4201 isis1                                           Dynamic
   1921.6800.4202 isis2                                           Static
   1921.6800.4203 isis3                                           Dynamic
```

show isis interface

Syntax
show isis interface
<brief | detail | extensive>
<*interface-name*>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show isis interface
<brief | detail | extensive>
<*interface-name*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display status information about Intermediate System-to-Intermediate System (IS-IS)-enabled interfaces.

Options
none—Display standard information about all IS-IS-enabled interfaces.

brief | detail | extensive—(Optional) Display the specified level of output.

interface-name—(Optional) Display information about the specified interface only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show isis interface on page 1853
show isis interface brief on page 1853
show isis interface detail on page 1853
show isis interface extensive on page 1853
show isis interface extensive (with LDP) on page 1854

Output Fields
Table 225 on page 1851 describes the output fields for the **show isis interface** command. Output fields are listed in the approximate order in which they appear.

Table 225: show isis interface Output Fields

Field Name	Field Description	Level of Output
interface-name	Name of the interface.	**detail**
Designated router	Routing device selected by other routers that is responsible for sending link-state advertisements that describe the network. Used only on broadcast networks.	**detail**
Index	Interface index assigned by the Junos OS kernel.	**detail**
State	Internal implementation information.	**detail**
Circuit id	Circuit identifier.	**detail**

Table 225: show isis interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
Circuit type	Circuit type: • **1**—Level 1 only • **2**—Level 2 only • **3**—Level 1 and Level 2	**detail**
LSP interval	Interval between link-state PDUs sent from the interface.	**detail**
CSNP interval	Interval between complete sequence number PDUs sent from the interface.	**detail extensive**
Sysid	System identifier.	**detail**
Interface	Interface through which the adjacency is made.	none **brief**
L or Level	Level: • **1**—Level 1 only • **2**—Level 2 only • **3**—Level 1 and Level 2	All levels
CirID	Circuit identifier.	none **brief**
Level 1 DR	Level 1 designated intermediate system.	none **brief**
Level 2 DR	Level 2 designated intermediate system.	none **brief**
L1/L2 Metric	Interface's metric for Level 1 and Level 2. If there is no information, the metric is **0**.	none **brief**
Adjacency advertisement: Advertise	This routing device has signaled not to advertise this interface to its neighbors in their label-switched paths (LSPs).	**detail extensive**
Adjacency advertisement: Suppress	This neighbor has signaled not to advertise this interface in the routing device's outbound LSPs.	**detail extensive**
Adjacencies	Number of adjacencies established on this interface.	**detail**
Priority	Priority value for this interface.	**detail**
Metric	Metric value for this interface.	**detail**
Hello(s) / Hello Interval	Interface's hello interval.	**detail extensive**
Hold(s) / Hold Time	Interface's hold time.	**detail extensive**

Table 225: show isis interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
Designated Router	Router responsible for sending network link-state advertisements, which describe all the routers attached to the network.	**detail**
Hello padding	Type of hello padding: • **Adaptive**—On point-to-point connections, the hello packets are padded from the initial detection of a new neighbor until the neighbor verifies the adjacency as Up in the adjacency state TLV. If the neighbor does not support the adjacency state TLV, then padding continues. On LAN connections, padding starts from the initial detection of a new neighbor until there is at least one active adjacency on the interface. • **Loose**—(Default) The hello packet is padded from the initial detection of a new neighbor until the adjacency transitions to the Up state. • **Strict**—Padding is performed on all interface types and for all adjacency states, and is continuous.	**extensive**
LDP sync state	Current LDP synchronization state: **in sync**, **in holddown**, or **not supported**.	**extensive**
reason	Reason for being in the LDP sync state.	**extensive**
config holdtime	Configured value of the hold timer.	**extensive**
remaining	If the state is not in sync and the hold time is not infinity, then this field displays the number of seconds remaining.	**extensive**

show isis interface

```
user@host> show isis interface
IS-IS interface database:
Interface             L CirID Level 1 DR       Level 2 DR        L1/L2 Metric
at-2/3/0.0            3   0x1 Point to Point   Point to Point         10/10
lo0.0                 0   0x1 Passive          Passive                 0/0
```

show isis interface brief

The output for the **show isis interface brief** command is identical to that for the **show isis interface** command. For sample output, see **show isis interface on page 1853**.

show isis interface detail

```
user@host> show isis interface detail
IS-IS interface database:
at-2/3/0.0
  Index: 66, State: 0x6, Circuit id: 0x1, Circuit type: 3
  LSP interval: 100 ms, CSNP interval: 5 s
  Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
    1             1       64     10     9.000       27
    2             1       64     10     9.000       27
lo0.0
  Index: 64, State: 0x6, Circuit id: 0x1, Circuit type: 0
  LSP interval: 100 ms, CSNP interval: disabled
  Level Adjacencies Priority Metric Hello (s) Hold (s) Designated Router
    1             0       64      0 Passive
    2             0       64      0 Passive
```

show isis interface extensive

```
user@host> show isis interface extensive
```

```
IS-IS interface database:
at-2/3/0.0
  Index: 66, State: 0x6, Circuit id: 0x1, Circuit type: 3
  LSP interval: 100 ms, CSNP interval: 5 s, Loose Hello padding
  Level 1
    Adjacencies: 1, Priority: 64, Metric: 10
    Hello Interval: 9.000 s, Hold Time: 27 s
  Level 2
    Adjacencies: 1, Priority: 64, Metric: 10
    Hello Interval: 9.000 s, Hold Time: 27 s
lo0.0
  Index: 64, State: 0x6, Circuit id: 0x1, Circuit type: 0
  LSP interval: 100 ms, CSNP interval: disabled, Loose Hello padding
  Level 1
    Adjacencies: 0, Priority: 64, Metric: 0
    Passive
  Level 2
    Adjacencies: 0, Priority: 64, Metric: 0
    Passive
```

show isis interface extensive (with LDP)

```
user@host> show isis interface extensive
IS-IS interface database:
so-1/1/2.0
  Index: 114, State: 0x6, Circuit id: 0x1, Circuit type: 2
  LSP interval: 100 ms, CSNP interval: 20 s, Loose Hello padding
  Adjacency advertisement: Advertise
  LDP sync state: in sync, for: 00:01:28, reason: LDP up during config
  config holdtime: 20 seconds
  Level 2
    Adjacencies: 1, Priority: 64, Metric: 11
    Hello Interval: 9.000 s, Hold Time: 27 s
    IPV4 MulticastMetric: 10
    IPV6 UnicastMetric: 10
```

show isis overview

Syntax **show isis overview**
<instance *instance-name*>
<**logical-system (all | *logical-system-name*)**>

Syntax (J-EX Series Switch) **show isis overview**
<instance *instance-name*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display Intermediate System-to Intermediate System (IS-IS) overview information.

Options none—Display standard overview information about IS-IS for all routing instances.

instance *instance-name*—(Optional) Display overview information for the specified routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level view

List of Sample Output **show isis overview on page 1856**

Output Fields Table 226 on page 1855 lists the output fields for the **show isis overview** command. Output fields are listed in the approximate order in which they appear.

Table 226: show isis overview Output Fields

Field Name	Field Description
instance	The IS-IS routing intance.
Router ID	Router ID of the routing device.
Adjacency holddown	Adjacency holddown capability: **enabled** or **disabled**.
Maximum Areas	Maximum number of IS-IS areas advertised by the routing device.
LSP life time	Lifetime of the link-state PDU, in seconds.
Attached bit evaluation	Attached bit capability: **enabled** or **disabled**.
SPF delay	Delay before performing consecutive Shortest Path First calculations.
SPF holddown	Delay before performing additional Shortest Path First (SPF) calculations after the maximum number of consecutive SPF calculations is reached.
SPF rapid runs	Maximum number of Shortest Path First calculations that can be performed in succession before the holddown timer begins.

Table 226: show isis overview Output Fields *(continued)*

Field Name	Field Description
Overload bit at startup is set	Overload bit capability is enabled.
Overload high metrics	Overload high metrics capability: **enabled** or **disabled**.
Overload timeout	Time period after which overload is reset and the time that remains before the timer is set to expire.
Traffic engineering	Traffic engineering capability: **enabled** or **disabled**.
Restart	Graceful restart capability: **enabled** or **disabled**.
Restart duration	Time period for complete reacquisition of IS-IS neighbors.
Helper mode	Graceful restart helper capability: **enalbed** or **disabled**.
Level	IS-IS level: • **1**—Level 1 information • **2**—Level 2 information
IPv4 is enabled	IP Protocol version 4 capability is enabled.
IPv6 is enabled	IP Protocol version 6 capability is enabled.
CLNS is enabled	OSI CLNP Protocol capability is enabled. (J Series routers only)
Internal route preference	Preference value of internal routes.
External route preference	Preference value of external routes.
Wide area metrics are enabled	Wide area metrics capability is enabled.
Narrow metrics is enabled	Narrow metrics capability is enabled.

show isis overview `user@host> show isis overview`

Sample Output

```
Instance: master
  Router ID: 192.168.1.220
  Adjacency holddown: enabled
  Maximum Areas: 3
  LSP life time: 65535
  Attached bit evaluation: enabled
  SPF delay: 200 msec, SPF holddown: 5000 msec, SPF rapid runs: 3
  Overload bit at startup is set
    Overload high metrics: disabled
```

```
  Overload timeout: 300 sec, expires in 295 seconds
IPv4 is enabled, IPv6 is enabled
Traffic engineering: enabled
Restart: Enabled
  Restart duration: 210 sec
  Helper mode: Enabled
Level 1
  Internal route preference: 15
  External route preference: 160
  Wide metrics are enabled, Narrow metrics are enabled
Level 2
  Internal route preference: 18
  External route preference: 165
  Wide metrics are enabled
```

show isis route

Syntax

show isis route
<*destination*>
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<topology (ipv4-multicast | ipv6-multicast | ipv6-unicast | unicast)>

Syntax (J-EX Series Switch)

show isis route
<*destination*>
<inet | inet6>
<instance *instance-name*>
<topology (ipv4-multicast | ipv6-multicast | ipv6-unicast | unicast)>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display the routes in the Intermediate System-to-Intermediate System (IS-IS) routing table.

Options

none—Display all routes in the IS-IS routing table for all supported address families for all routing instances.

destination—(Optional) Destination address for the route.

inet | inet6—(Optional) Display inet (IPv4) or inet6 (IPv6) routes, respectively.

instance *instance-name*—(Optional) Display routes for the specified routing instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

topology (ipv4-multicast | ipv6-multicast | ipv6-unicast | unicast)—(Optional) Display routes for the specified topology only, or use unicast to display information, if available, for both IPv4 and IPv6 unicast topologies.

Required Privilege Level

view

List of Sample Output

show isis route logical-system on page 1859
show isis route (CLNS) on page 1860

Output Fields

Table 227 on page 1859 describes the output fields for the **show isis route** command. Output fields are listed in the approximate order in which they appear.

Table 227: show isis route Output Fields

Field Name	Field Description
Current version	Number of the current version of the IS-IS routing table.
L1	Version of Level 1 SPF that was run.
L2	Version of Level 2 SPF that was run.
Prefix	Destination of the route.
L	IS-IS level: • **1**—Level 1 only • **2**—Level 2 only • **3**—Level 1 and Level 2
Version	Version of SPF that generated the route.
Metric	Metric value associated with the route.
Type	Metric type: **int** (internal) or **ext** (external).
Interface	Interface to the next hop.
Via	System identifier of the next hop, displayed as a name if possible.
ISO Routes	ISO routing table entries.
snpa	MAC address.

show isis route logical-system

```
user@host> show isis route logical-system ls1
IS-IS routing table            Current version: L1: 8 L2: 11
Prefix             L Version Metric Type Interface    Via
10.9.7.0/30        2      11     20 int  gr-0/2/0.0   h
10.9.201.1/32      2      11     60 int  gr-0/2/0.0   h
IPV6 Unicast IS-IS routing table            Current version: L1: 9 L2: 11
Prefix             L Version Metric Type Interface    Via
8009:3::a09:3200/126 2      11     20 int gr-0/2/0.0  h
```

show isis route (CLNS)

```
user@host> show isis route
IS-IS routing table            Current version: L1: 10 L2: 8
IPv4/IPv6 Routes
Prefix             L Version   Metric Type Interface    Via
0.0.0.0/0          1      10       10 int  fe-0/0/1.0   ISIS.0
ISO Routes
Prefix  L   Version   Metric Type Interface    Via    snpa
0/0
        1        10       10 int  fe-0/0/1.0   isis.0 0:12:0:34:0:56
47.0005.80ff.f800.0000.0108.0001/104
        1        10        0 int
47.0005.80ff.f800.0000.0108.0001.1921.6800.4001/152
        1        10       10 int  fe-0/0/1.0 isis.0 0:12:0:34:0:56
47.0005.80ff.f800.0000.0108.0001.1921.6800.4002/152
        1        10       20 int  fe-0/0/1.0 isis.0 0:12:0:34:0:56
47.0005.80ff.f800.0000.0108.0002/104
        1        10        0 int
47.0005.80ff.f800.0000.0108.0002.1921.6800.4001/152
        1        10       10 int  fe-0/0/1.0 isis.0 0:12:0:34:0:56
```

show isis spf

Syntax

show isis spf (brief | log | results)
<instance *instance-name*>
<level (1 | 2)>
<logical-system (all | *logical-system-name*)>
<topology (ipv4-multicast | ipv6-multicast | ipv6-unicast | unicast)>

Syntax (J-EX Series Switch)

show isis spf (brief | log | results)
<instance *instance-name*>
<level (1 | 2)>
<topology (ipv4-multicast | ipv6-multicast | ipv6-unicast | unicast)>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display information about Intermediate System-to-Intermediate System (IS-IS) shortest-path-first (SPF) calculations.

Options

brief—Display an overview of SPF calculations.

log—Display the log of SPF calculations.

results—Display the results of SPF calculations.

instance *instance instance-name*—(Optional) Display SPF calculations for the specified routing instance.

level (1 | 2)—(Optional) Display SPF calculations for the specified IS-IS level.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

topology (ipv4-multicast | ipv6-multicast | ipv6-unicast | unicast)—(Optional) Display SPF calculations for the specified topology only.

Required Privilege Level

view

List of Sample Output

show isis spf brief on page 1862
show isis spf log on page 1863
show isis spf results on page 1864
show isis spf results (CLNS) on page 1865

Output Fields

Table 228 on page 1861 describes the output fields for the **show isis spf** command. Output fields are listed in the approximate order in which they appear.

Table 228: show isis spf Output Fields

Field Name	Field Description
Node	System ID of a node.
Metric	Metric to the node.

Table 228: show isis spf Output Fields *(continued)*

Field Name	Field Description
Interface	Interface of the next hop.
Via	System ID of the next hop.
SNPA	Subnetwork point of attachment (MAC address of the next hop).
Start time	(**log** option only) Time that the SPF computation started.
Elapsed (secs)	(**log** option only) Length of time, in seconds, required to complete the SPF computation.
Count	(**log** option only) Number of times the SPF was triggered.
Reason	(**log** option only) Reason that the SPF computation was completed.

show isis spf brief

```
user@host> show isis spf brief logical-system ls1
 IS-IS level 1 SPF results:
Node             Metric     Interface         Via             SNPA
scat.00          10         ge-1/1/0.0        scat            0:90:69:a6:48:9d
fix.02           10
fix.00           0
  3 nodes

 IS-IS level 2 SPF results:
Node             Metric     Interface         Via             SNPA
skag.00          20         gr-0/2/0.0        h
skag.02          20         gr-0/2/0.0        h
h.00             10         gr-0/2/0.0        h
fix.00           0
  4 nodes

IPV6 Unicast IS-IS level 1 SPF results:
Node             Metric     Interface         Via             SNPA
scat.00          10         ge-1/1/0.0        scat            0:90:69:a6:48:9d
                            ge-1/1/0.0        scat            0:90:69:a6:48:9d
fix.02           10
fix.00           0
  3 nodes

IPV6 Unicast IS-IS level 2 SPF results:
Node             Metric     Interface         Via             SNPA
skag.00          20         gr-0/2/0.0        h
                            gr-0/2/0.0        h
skag.02          20         gr-0/2/0.0        h
                            gr-0/2/0.0        h
h.00             10         gr-0/2/0.0        h
                            gr-0/2/0.0        h
fix.00           0
  4 nodes

Multicast IS-IS level 1 SPF results:
Node             Metric     Interface         Via             SNPA
scat.00          10         ge-1/1/0.0        scat            0:90:69:a6:48:9d
```

```
fix.02           10
fix.00           0
  3 nodes

Multicast IS-IS level 2 SPF results:
Node             Metric     Interface         Via             SNPA
skag.00          20         gr-0/2/0.0       h
skag.02          20         gr-0/2/0.0       h
h.00             10         gr-0/2/0.0       h
fix.00           0
  4 nodes
```

show isis spf log

```
user@host> show isis spf log logical-system lsl
 IS-IS level 1 SPF log:
Start time          Elapsed (secs) Count Reason
Fri Oct 31 12:41:18        0.000069    1 Reconfig
Fri Oct 31 12:41:18        0.000107    3 Updated LSP fix.00-00
Fri Oct 31 12:41:18        0.000050    3 Address change on so-1/2/2.0
Fri Oct 31 12:41:23        0.000033    1 Updated LSP fix.00-00
Fri Oct 31 12:41:28        0.000178    5 New adjacency scat on ge-1/1/0.0
Fri Oct 31 12:41:59        0.000060    1 Updated LSP fix.00-00
Fri Oct 31 12:42:30        0.000161    2 Multi area attachment change
Fri Oct 31 12:56:58        0.000198    1 Periodic SPF
Fri Oct 31 13:10:29        0.000209    1 Periodic SPF
 IS-IS level 2 SPF log:

Start time          Elapsed (secs) Count Reason
Fri Oct 31 12:41:18        0.000035    1 Reconfig
Fri Oct 31 12:41:18        0.000047    2 Updated LSP fix.00-00
Fri Oct 31 12:41:18        0.000043    5 Address change on gr-0/2/0.0
Fri Oct 31 12:41:23        0.000022    1 Updated LSP fix.00-00
Fri Oct 31 12:41:59        0.000144    3 New adjacency h on gr-0/2/0.0
Fri Oct 31 12:42:30        0.000257    3 New LSP skag.00-00
Fri Oct 31 12:54:37        0.000195    1 Periodic SPF
Fri Oct 31 12:55:50        0.000178    1 Updated LSP fix.00-00
Fri Oct 31 12:55:55        0.000174    1 Updated LSP h.00-00
Fri Oct 31 12:55:58        0.000176    1 Updated LSP skag.00-00
Fri Oct 31 13:08:14        0.000198    1 Periodic SPF
IPV6 Unicast IS-IS level 1 SPF log:

Start time          Elapsed (secs) Count Reason
Fri Oct 31 12:41:18        0.000028    1 Reconfig
Fri Oct 31 12:41:18        0.000043    3 Updated LSP fix.00-00
Fri Oct 31 12:41:18        0.000112    4 Updated LSP fix.00-00
Fri Oct 31 12:41:23        0.000059    1 Updated LSP fix.00-00
Fri Oct 31 12:41:25        0.000041    1 Updated LSP fix.00-00
Fri Oct 31 12:41:28        0.000103    5 New adjacency scat on ge-1/1/0.0
Fri Oct 31 12:41:59        0.000040    1 Updated LSP fix.00-00
Fri Oct 31 12:42:30        0.000118    2 Multi area attachment change
Fri Oct 31 12:56:08        0.000289    1 Periodic SPF
Fri Oct 31 13:11:07        0.000214    1 Periodic SPF
IPV6 Unicast IS-IS level 2 SPF log:

Start time          Elapsed (secs) Count Reason
Fri Oct 31 12:41:18        0.000027    1 Reconfig
Fri Oct 31 12:41:18        0.000039    2 Updated LSP fix.00-00
Fri Oct 31 12:41:18        0.000049    6 Updated LSP fix.00-00
Fri Oct 31 12:41:23        0.000025    1 Updated LSP fix.00-00
Fri Oct 31 12:41:25        0.000023    1 Updated LSP fix.00-00
Fri Oct 31 12:41:59        0.000087    3 New adjacency h on gr-0/2/0.0
Fri Oct 31 12:42:30        0.000123    3 New LSP skag.00-00
```

```
Fri Oct 31 12:55:50        0.000121    1 Updated LSP fix.00-00
Fri Oct 31 12:55:55        0.000121    1 Updated LSP h.00-00
Fri Oct 31 12:55:58        0.000121    1 Updated LSP skag.00-00
Fri Oct 31 13:09:46        0.000201    1 Periodic SPF
...
```

show isis spf results

```
user@host> show isis spf results logical-system ls1
IS-IS level 1 SPF results:
Node             Metric     Interface         Via             SNPA
scat.00          10         ge-1/1/0.0        scat            0:90:69:a6:48:9d
                 20         10.9.1.0/30
fix.02           10
fix.00           0
                 10         10.9.1.0/30
                 10         10.9.5.0/30
                 10         10.9.6.0/30
                 20         10.9.7.0/30
                 60         10.9.201.1/32
  3 nodes

 IS-IS level 2 SPF results:
Node             Metric     Interface         Via             SNPA
skag.00          20         gr-0/2/0.0        h
                 30         10.9.7.0/30
skag.02          20         gr-0/2/0.0        h
h.00             10         gr-0/2/0.0        h
                 20         10.9.6.0/30
                 20         10.9.7.0/30
                 60         10.9.201.1/32
fix.00           0
                 10         10.9.1.0/30
                 10         10.9.5.0/30
                 10         10.9.6.0/30
  4 nodes

IPV6 Unicast IS-IS level 1 SPF results:
Node             Metric     Interface         Via             SNPA
scat.00          10         ge-1/1/0.0        scat            0:90:69:a6:48:9d
                            ge-1/1/0.0        scat            0:90:69:a6:48:9d
                 20         8009:1::a09:1400/126
fix.02           10
fix.00           0
                 10         8009:1::a09:1400/126
                 10         8009:2::a09:1e00/126
                 20         8009:3::a09:3200/126
                 10         8009:4::a09:2800/126
  3 nodes

IPV6 Unicast IS-IS level 2 SPF results:
Node             Metric     Interface         Via             SNPA
skag.00          20         gr-0/2/0.0        h
                            gr-0/2/0.0        h
                 30         8009:3::a09:3200/126
skag.02          20         gr-0/2/0.0        h
                            gr-0/2/0.0        h
h.00             10         gr-0/2/0.0        h
                            gr-0/2/0.0        h
                 20         8009:3::a09:3200/126
                 20         8009:4::a09:2800/126
fix.00           0
                 10         8009:1::a09:1400/126
```

```
                10          8009:2::a09:1e00/126
                10          8009:4::a09:2800/126
  4 nodes

Multicast IS-IS level 1 SPF results:
Node            Metric     Interface         Via             SNPA
scat.00         10         ge-1/1/0.0        scat            0:90:69:a6:48:9d
fix.02          10
fix.00          0
  3 nodes

Multicast IS-IS level 2 SPF results:
Node            Metric     Interface         Via             SNPA
skag.00         20         gr-0/2/0.0        h
skag.02         20         gr-0/2/0.0        h
h.00            10         gr-0/2/0.0        h
fix.00          0
  4 nodes
...
```

show isis spf results (CLNS)

```
user@host> show isis spf results
IS-IS level 1 SPF results:
Node            Metric     Interface         Via             SNPA
skag.00 10                 fe-0/0/1.0        toothache       0:12:0:34:0:56
                           fe-0/0/1.0        toothache       0:12:0:34:0:56
                20         192.168.37.64/29
                10         1921.6800.4001
                20         1921.6800.4002
pro1-a.02       10
pro1-a.00       0
                0          10.255.245.1/32
                10         192.168.37.64/29
                0          1921.6800.4211
  3 nodes

IS-IS level 2 SPF results:
Node            Metric     Interface         Via             SNPA
skag.00 10                 fe-0/0/1.0        toothache       0:12:0:34:0:56
                           fe-0/0/1.0        toothache       0:12:0:34:0:56
                20         10.255.245.1/32
                20         192.168.37.64/29
                20         47.0005.80ff.f800.0000.0109.0010/104
pro1-a.02       10
pro1-a.00       0
                0          10.255.245.1/32
                10         192.168.37.64/29
  3 nodes
```

show isis statistics

Syntax
show isis statistics
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show isis statistics
<instance *instance-name*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display statistics about Intermediate System-to-Intermediate System (IS-IS) traffic.

Options
none—Display IS-IS traffic statistics for all routing instances.

instance *instance-name*—(Optional) Display statistics for the specified routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- **clear isis statistics on page 1766**

List of Sample Output
show isis statistics on page 1867

Output Fields
Table 229 on page 1866 describes the output fields for the **show isis statistics** command. Output fields are listed in the approximate order in which they appear.

Table 229: show isis statistics Output Fields

Field Name	Field Description
PDU type	Protocol data unit type: • **CSNP**—Complete sequence number PDUs contain a complete list of all link-state PDUs in the IS-IS database. CSNPs are sent periodically on all links, and the receiving systems use the information in the CSNP to update and synchronize their link-state PDU databases. The designated router multicasts CSNPs on broadcast links in place of sending explicit acknowledgments for each link-state PDU. • **IIH**—IS-IS hello packets are broadcast to discover the identity of neighboring IS-IS systems and to determine whether the neighbors are Level 1 or Level 2 intermediate systems. • **LSP**—Link-state PDUs contain information about the state of adjacencies to neighboring IS-IS systems. Link-state PDUs are flooded periodically throughout an area. • **PSNP**—Partial sequence number PDUs are sent multicast by a receiver when it detects that it is missing a link-state PDU; that is, when its link-state PDU database is out of date. The receiver sends a PSNP to the system that transmitted the CSNP, effectively requesting that the missing link-state PDU be transmitted. That routing device, in turn, forwards the missing link-state PDU to the requesting routing device. • **Unknown**—The PDU type is unknown.
Received	Number of PDUs received since IS-IS started or since the statistics were set to zero.

Table 229: show isis statistics Output Fields *(continued)*

Field Name	Field Description
Processed	Number of PDUs received less the number dropped.
Drops	Number of PDUs dropped.
Sent	Number of PDUs transmitted since IS-IS started or since the statistics were set to zero.
Rexmit	Number of PDUs retransmitted since IS-IS started or since the statistics were set to zero.
Total packets received/sent	Total number of PDUs received and transmitted since IS-IS started or since the statistics were set to zero.
SNP queue length	Number of CSPN and PSNP packets currently waiting in the queue for processing. This value is almost always **0**.
LSP queue length	Number of link-state PDUs waiting in the queue for processing. This value is almost always **0**.
SPF runs	Number of shortest-path-first (SPF) calculations that have been performed. If this number is incrementing rapidly, it indicates that the network is unstable.
Fragments rebuilt	Number of link-state link-state PDU fragments that the local system has computed.
LSP regenerations	Number of link-state PDUs that have been regenerated. A link state PDU is regenerated when it is nearing the end of its lifetime and it has not changed.
Purges initiated	Number of purges that the system initiated. A purge is initiated if the software decides that a link-state PDU must be removed from the network.

show isis statistics

```
user@host> show isis statistics
IS-IS statistics for merino:

PDU type     Received  Processed      Drops       Sent     Rexmit
LSP             12227      12227          0       8184        683
IIH            113808     113808          0     115817          0
CSNP           198868     198868          0     198934          0
PSNP             6985       6979          6       8274          0
Unknown             0          0          0          0          0
Totals         331888     331882          6     331209        683

Total packets received: 331888 Sent: 331892

SNP queue length:           0 Drops:          0
LSP queue length:           0 Drops:          0

SPF runs:                1014
Fragments rebuilt:       1038
LSP regenerations:        425
Purges initiated:           0
```

show ospf3 database

Syntax

show ospf3 database
<brief | detail | extensive | summary>
<advertising-router (*address* | self)>
<area *area-id*>
<external>
<instance *instance-name*>
<inter-area-prefix>
<inter-area-router>
<intra-area-prefix>
<link>
<link-local>
<logical-system (all | *logical-system-name*)>
<lsa-id *lsa-id*>
<network>
<nssa>
<realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)>
<router>

Syntax (J-EX Series Switch)

show ospf3 database
<brief | detail | extensive | summary>
<advertising-router (*address* | self)>
<area *area-id*>
<external>
<instance *instance-name*>
<inter-area-prefix>
<inter-area-router>
<intra-area-prefix>
<link>
<link-local>
<lsa-id *lsa-id*>
<network>
<nssa>
<router>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display the entries in the Open Shortest Path First version 3 (OSPFv3) link-state database, which contains data about link-state advertisement (LSA) packets.

Options

none—Display standard information about all entries in the OSPFv3 link-state database.

brief | detail | extensive | summary—(Optional) Display the specified level of output.

advertising-router (*address* | self)—(Optional) Display the LSAs advertised either by a particular routing device or by this routing device.

area *area-id*—(Optional) Display the LSAs in a particular area.

external—(Optional) Display external LSAs.

instance *instance-name*—(Optional) Display all OSPF database information under the named routing instance.

inter-area-prefix—(Optional) Display information about interarea-prefix LSAs.

inter-area-router—(Optional) Display information about interarea-router LSAs.

intra-area-prefix—(Optional) Display information about intra-area-prefix LSAs.

link—(Optional) Display information about link LSAs.

link-local—(Optional) Display information about link-local LSAs.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

lsa-id *lsa-id*—(Optional) Display the LSA with the specified LSA identifier.

network—(Optional) Display information about network LSAs.

nssa—(Optional) Display information about not-so-stubby area (NSSA) LSAs.

realm (ipv4-multicast | ipv4-unicast | ipv6-multicast)—(Optional) Display information about the specified OSPFv3 realm, or address family. Use the **realm** option to specify an address family other than IPv6 unicast, which is the default.

router—(Optional) Display information about router LSAs.

Required Privilege Level view

Related Documentation
- **clear (ospf | ospf3) database on page 1748**

List of Sample Output
show ospf3 database brief on page 1874
show ospf3 database extensive on page 1874
show ospf3 database summary on page 1877

Output Fields Table 230 on page 1869 lists the output fields for the **show ospf3 database** command. Output fields are listed in the approximate order in which they appear.

Table 230: show ospf3 database Output Fields

Field Name	Field Description	Level of Output
OSPF link state database, area *area-number*	Entries in the link-state database for this area.	**brief detail extensive**
OSPF AS SCOPE link state database	Entries in the AS scope link-state database.	**brief detail extensive**

Table 230: show ospf3 database Output Fields *(continued)*

Field Name	Field Description	Level of Output
OSPF Link-Local link state database, interface *interface-name*	Entries in the link-local link-state database for this interface.	**brief detail extensive**
area	Area number. Area **0.0.0.0** is the backbone area.	All levels
Type	Type of link advertisement: **Extern**, **InterArPfx**, **InterArRtr**, **IntraArPrx** , **Link**, **Network**, **NSSA**, or **Router**.	**brief detail extensive**
ID	Link identifier included in the advertisement. An asterisk (*) preceding the identifier marks database entries that originated from the local routing device.	**brief detail extensive**
Adv Rtr	Address of the routing device that sent the advertisement.	**brief detail extensive**
Seq	Link sequence number of the advertisement.	**brief detail extensive**
Age	Time elapsed since the LSA was originated, in seconds.	**brief detail extensive**
Cksum	Checksum value of the LSA.	**brief detail extensive**
Len	Length of the advertisement, in bytes.	**brief detail extensive**
Router (Router Link-State Advertisements)		
bits	Flags describing the routing device that generated the LSP.	**detail extensive**
Options	Option bits carried in the router LSA.	**detail extensive**
For Each Router Link		
Type	Type of interface. The value of all other output fields describing a routing device interface depends on the interface's type: • **PointToPoint (1)**—Point-to-point connection to another routing device. • **Transit (2)**—Connection to a transit network. • **Virtual (4)**—Virtual link.	**detail extensive**
Loc-if-id	Local interface ID assigned to the interface that uniquely identifies the interface with the routing device.	**detail extensive**
Nbr-if-id	Interface ID of the neighbor's interface for this routing device link.	**detail extensive**
Nbr-rtr-id	Router ID of the neighbor routing device (for type 2 interfaces, the attached link's designated router).	**detail extensive**
Metric	Cost of the router link.	**detail extensive**
Gen timer	How long until the LSA is regenerated, in the format ***hours*:*minutes*:*seconds***.	**extensive**

Table 230: show ospf3 database Output Fields *(continued)*

Field Name	Field Description	Level of Output
Aging timer	How long until the LSA expires, in the format ***hours*:*minutes*:*seconds***.	**extensive**
Installed *nn*:*nn*:*nn* ago	How long ago the route was installed, in the format ***hours*:*minutes*:*seconds***.	**extensive**
expires in *nn*:*nn*:*nn*	How long until the route expires, in the format ***hours*:*minutes*:*seconds***.	**extensive**
sent *nn*:*nn*:*nn* ago	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format ***hours*:*minutes*:*seconds***.	**extensive**
Ours	Indicates that this is a local advertisement.	**extensive**
Network (Network Link-State Advertisements)		
Options	Option bits carried in the network LSA.	**detail extensive**
Attached Router	Router IDs of each of the routing devices attached to the link. Only routing devices that are fully adjacent to the designated router are listed. The designated router includes itself in this list.	**detail extensive**
InterArPfx (Interarea-Prefix Link-State Advertisements)		
Prefix	IPv6 address prefix.	**detail extensive**
Prefix-options	Option bit associated with the prefix.	**detail extensive**
Metric	Cost of this route. Expressed in the same units as the interface costs in the router LSAs. When the interarea-prefix LSA is describing a route to a range of addresses, the cost is set to the maximum cost to any reachable component of the address range.	**detail extensive**
Gen timer	How long until the LSA is regenerated, in the format ***hours*:*minutes*:*seconds***.	**extensive**
Aging timer	How long until the LSA expires, in the format ***hours*:*minutes*:*seconds***.	**extensive**
Installed *nn*:*nn*:*nn* ago	How long ago the route was installed, in the format ***hours*:*minutes*:*seconds***.	**extensive**
expires in *nn*:*nn*:*nn*	How long until the route expires, in the format ***hours*:*minutes*:*seconds***.	**extensive**
sent *nn*:*nn*:*nn* ago	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format ***hours*:*minutes*:*seconds***.	**extensive**
Ours	Indicates that this is a local advertisement.	**extensive**
InterArRtr (Interarea-Router Link-State Advertisements)		
Dest-router-id	Router ID of the routing device described by the LSA.	**detail extensive**
options	Optional capabilities supported by the routing device.	**detail extensive**

Table 230: show ospf3 database Output Fields *(continued)*

Field Name	Field Description	Level of Output
Metric	Cost of this route. Expressed in the same units as the interface costs in the router LSAs. When the interarea-prefix LSA is describing a route to a range of addresses, the cost is set to the maximum cost to any reachable component of the address range.	**detail extensive**
Prefix	IPv6 address prefix.	**extensive**
Prefix-options	Option bit associated with the prefix.	**extensive**
Extern (External Link-State Advertisements)		
Prefix	IPv6 address prefix.	**detail extensive**
Prefix-options	Option bit associated with the prefix.	**detail extensive**
Metric	Cost of the route, which depends on the value of **Type**.	**detail extensive**
Type *n*	Type of external metric: **Type 1** or **Type 2**.	**detail extensive**
Aging timer	How long until the LSA expires, in the format ***hours*:*minutes*:*seconds***.	**extensive**
Installed *nn:nn:nn* **ago**	How long ago the route was installed, in the format ***hours*:*minutes*:*seconds***.	**extensive**
expires in *nn:nn:nn*	How long until the route expires, in the format ***hours*:*minutes*:*seconds***.	**extensive**
sent *nn:nn:nn* **ago**	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format ***hours*:*minutes*:*seconds***.	**extensive**
Link (Link-State Advertisements)		
IPv6-Address	IPv6 link-local address on the link for which this link LSA originated.	**detail extensive**
Options	Option bits carried in the link LSA.	**detail extensive**
priority	Router priority of the interface attaching the originating routing device to the link.	**detail extensive**
Prefix-count	Number of IPv6 address prefixes contained in the LSA. The rest of the link LSA contains a list of IPv6 prefixes to be associated with the link.	**detail extensive**
Prefix	IPv6 address prefix.	**detail extensive**
Prefix-options	Option bit associated with the prefix.	**detail extensive**
Gen timer	How long until the LSA is regenerated, in the format ***hours*:*minutes*:*seconds***.	**extensive**
Aging timer	How long until the LSA expires, in the format ***hours*:*minutes*:*seconds***.	**extensive**

Table 230: show ospf3 database Output Fields *(continued)*

Field Name	Field Description	Level of Output
Installed *nn:nn:nn* ago	How long ago the route was installed, in the format ***hours*:*minutes*:*seconds***.	**extensive**
expires in *nn:nn:nn*	How long until the route expires, in the format ***hours*:*minutes*:*seconds***.	**extensive**
sent *nn:nn:nn* ago	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format ***hours*:*minutes*:*seconds***.	**extensive**
Ours	Indicates that this is a local advertisement.	**extensive**
IntraArPfx (Intra-Area-Prefix Link-State Advertisements)		
Ref-lsa-type	LSA type of the referenced LSA. • **Router**—Address prefixes are associated with a router LSA. • **Network**—Address prefixes are associated with a network LSA.	**detail extensive**
Ref-lsa-id	Link-state ID of the referenced LSA.	**detail extensive**
Ref-router-id	Advertising router ID of the referenced LSA.	**detail extensive**
Prefix-count	Number of IPv6 address prefixes contained in the LSA. The rest of the link LSA contains a list of IPv6 prefixes to be associated with the link.	**detail extensive**
Prefix	IPv6 address prefix.	**detail extensive**
Prefix-options	Option bit associated with the prefix.	**detail extensive**
Metric	Cost of this prefix. Expressed in the same units as the interface costs in the router LSAs.	**detail extensive**
Gen timer	How long until the LSA is regenerated, in the format ***hours*:*minutes*:*seconds***.	**extensive**
Aging timer	How long until the LSA expires, in the format ***hours*:*minutes*:*seconds***.	**extensive**
Installed *hh*:*mm*:*ss* ago	How long ago the route was installed, in the format ***hours*:*minutes*:*seconds***.	**extensive**
expires in *hh*:*mm*:*ss*	How long until the route expires, in the format ***hours*:*minutes*:*seconds***.	**extensive**
sent *hh*:*mm*:*ss* ago	Time elapsed since the LSA was last transmitted or flooded to an adjacency or an interface, respectively, in the format ***hours*:*minutes*:*seconds***.	**extensive**
***n* Router LSAs**	Number of router LSAs in the link-state database.	**summary**
***n* Network LSAs**	Number of network LSAs in the link-state database.	**summary**
***n* InterArPfx LSAs**	Number of interarea-prefix LSAs in the link-state database.	**summary**

Table 230: show ospf3 database Output Fields *(continued)*

Field Name	Field Description	Level of Output
***n* InterArRtr LSAs**	Number of interarea-router LSAs in the link-state database.	**summary**
***n* IntraArPfx LSAs**	Number of intra-area-prefix LSAs in the link-state database.	**summary**
Externals	Display of the external LSA database.	**summary**
***n* Extern LSAs**	Number of external LSAs in the link-state database.	**summary**
Interface ***interface-name***	Name of the interface for which link-local LSA information is displayed.	**summary**
***n* Link LSAs**	Number of link LSAs in the link-state database.	**summary**

show ospf3 database brief

```
user@host> show ospf3 database brief
    OSPF3 link state database, area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Router      0.0.0.1          10.255.4.85       0x80000003   885  0xa697  40
Router     *0.0.0.1          10.255.4.93       0x80000002   953  0xc677  40
InterArPfx *0.0.0.2          10.255.4.93       0x80000001   910  0xb96f  44
InterArRtr *0.0.0.1          10.255.4.93       0x80000001   910  0xe159  32
IntraArPfx *0.0.0.1          10.255.4.93       0x80000002   432  0x788f  72

    OSPF3 link state database, area 0.0.0.1
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Router     *0.0.0.1          10.255.4.93       0x80000003   916  0xea40  40
Router      0.0.0.1          10.255.4.97       0x80000006   851  0xc95b  40
Network     0.0.0.2          10.255.4.97       0x80000002   916  0x4598  32
InterArPfx *0.0.0.1          10.255.4.93       0x80000002   117  0xa980  44
InterArPfx *0.0.0.2          10.255.4.93       0x80000002    62  0xd47e  44
NSSA        0.0.0.1          10.255.4.97       0x80000002   362  0x45ee  44
IntraArPfx  0.0.0.1          10.255.4.97       0x80000006   851  0x2f77  52

    OSPF3 AS SCOPE link state database
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Extern      0.0.0.1          10.255.4.85       0x80000002    63  0x9b86  44
Extern     *0.0.0.1          10.255.4.93       0x80000001   910  0x59c9  44

    OSPF3 Link-Local link state database, interface ge-1/3/0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link       *0.0.0.2          10.255.4.93       0x80000003   916  0x4dab  64
```

show ospf3 database extensive

```
user@host> show ospf3 database extensive
    OSPF3 link state database, area 0.0.0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Router      0.0.0.1          10.255.4.85       0x80000003  1028  0xa697  40
  bits 0x2, Options 0x13
  Type PointToPoint (1), Metric 10
      Loc-If-Id 2, Nbr-If-Id 3, Nbr-Rtr-Id 10.255.4.93
  Aging timer 00:42:51
  Installed 00:17:05 ago, expires in 00:42:52, sent 02:37:54 ago
Router     *0.0.0.1          10.255.4.93       0x80000002  1096  0xc677  40
  bits 0x3, Options 0x13
  Type PointToPoint (1), Metric 10
```

```
      Loc-If-Id 3, Nbr-If-Id 2, Nbr-Rtr-Id 10.255.4.85
  Gen timer 00:00:40
  Aging timer 00:41:44
  Installed 00:18:16 ago, expires in 00:41:44, sent 00:18:14 ago
  Ours
InterArPfx *0.0.0.2          10.255.4.93      0x80000001  1053  0xb96f  44
  Prefix feee::10:10:2:0/126
  Prefix-options 0x0, Metric 10
  Gen timer 00:17:02
  Aging timer 00:42:26
  Installed 00:17:33 ago, expires in 00:42:27, sent 00:17:31 ago
  Ours
InterArPfx *0.0.0.3          10.255.4.93      0x80000001  1053  0x71d3  44
  Prefix feee::10:255:4:97/128
  Prefix-options 0x0, Metric 10
  Gen timer 00:21:07
  Aging timer 00:42:26
  Installed 00:17:33 ago, expires in 00:42:27, sent 00:17:31 ago
  Ours
InterArRtr *0.0.0.1          10.255.4.93      0x80000001  1053  0xe159  32
  Dest-router-id 10.255.4.97, Options 0x19, Metric 10
  Gen timer 00:29:18
  Aging timer 00:42:26
  Installed 00:17:33 ago, expires in 00:42:27, sent 00:17:31 ago
  Ours
IntraArPfx  0.0.0.1          10.255.4.85      0x80000002  1028  0x2403  72
  Ref-lsa-type Router, Ref-lsa-id 0.0.0.0, Ref-router-id 10.255.4.85
  Prefix-count 2
  Prefix feee::10:255:4:85/128
    Prefix-options 0x2, Metric 0
  Prefix feee::10:10:1:0/126
    Prefix-options 0x0, Metric 10
  Aging timer 00:42:51
  Installed 00:17:05 ago, expires in 00:42:52, sent 02:37:54 ago
IntraArPfx *0.0.0.1          10.255.4.93      0x80000002   575  0x788f  72
  Ref-lsa-type Router, Ref-lsa-id 0.0.0.0, Ref-router-id 10.255.4.93
  Prefix-count 2
  Prefix feee::10:255:4:93/128
    Prefix-options 0x2, Metric 0
  Prefix feee::10:10:1:0/126
    Prefix-options 0x0, Metric 10
  Gen timer 00:33:23
  Aging timer 00:50:24
  Installed 00:09:35 ago, expires in 00:50:25, sent 00:09:33 ago
    OSPF3 link state database, area 0.0.0.1
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Router     *0.0.0.1          10.255.4.93      0x80000003  1059  0xea40  40
  bits 0x3, Options 0x19
  Type Transit (2), Metric 10
      Loc-If-Id 2, Nbr-If-Id 2, Nbr-Rtr-Id 10.255.4.97
  Gen timer 00:08:51
  Aging timer 00:42:20
  Installed 00:17:39 ago, expires in 00:42:21, sent 00:17:37 ago
Router      0.0.0.1          10.255.4.97      0x80000006   994  0xc95b  40
  bits 0x2, Options 0x19
  Type Transit (2), Metric 10
      Loc-If-Id 2, Nbr-If-Id 2, Nbr-Rtr-Id 10.255.4.97
  Aging timer 00:43:25
  Installed 00:16:31 ago, expires in 00:43:26, sent 02:37:54 ago
Network     0.0.0.2          10.255.4.97      0x80000002  1059  0x4598  32
  Options 0x11
```

```
  Attached router 10.255.4.97
  Attached router 10.255.4.93
  Aging timer 00:42:20
  Installed 00:17:36 ago, expires in 00:42:21, sent 02:37:54 ago
InterArPfx *0.0.0.1          10.255.4.93      0x80000002   260  0xa980  44
  Prefix feee::10:10:1:0/126
  Prefix-options 0x0, Metric 10
  Gen timer 00:45:39
  Aging timer 00:55:39
  Installed 00:04:20 ago, expires in 00:55:40, sent 00:04:18 ago
  Ours
InterArPfx *0.0.0.2          10.255.4.93      0x80000002   205  0xd47e  44
  Prefix feee::10:255:4:93/128
  Prefix-options 0x0, Metric 0
  Gen timer 00:46:35
  Aging timer 00:56:35
  Installed 00:03:25 ago, expires in 00:56:35, sent 00:03:23 ago
  Ours
InterArPfx *0.0.0.3          10.255.4.93      0x80000001  1089  0x9bbb  44
  Prefix feee::10:255:4:85/128
  Prefix-options 0x0, Metric 10
  Gen timer 00:04:46
  Aging timer 00:41:51
  Installed 00:18:09 ago, expires in 00:41:51, sent 00:17:43 ago
  Ours
NSSA        0.0.0.1          10.255.4.97      0x80000002   505  0x45ee  44
  Prefix feee::200:200:1:0/124
  Prefix-options 0x8, Metric 10, Type 2,
  Aging timer 00:51:35
  Installed 00:08:22 ago, expires in 00:51:35, sent 02:37:54 ago
IntraArPfx  0.0.0.1          10.255.4.97      0x80000006   994  0x2f77  52
  Ref-lsa-type Router, Ref-lsa-id 0.0.0.0, Ref-router-id 10.255.4.97
  Prefix-count 1
  Prefix feee::10:255:4:97/128
    Prefix-options 0x2, Metric 0
  Aging timer 00:43:25
  Installed 00:16:31 ago, expires in 00:43:26, sent 02:37:54 ago
IntraArPfx  0.0.0.3          10.255.4.97      0x80000002  1059  0x4446  52
  Ref-lsa-type Network, Ref-lsa-id 0.0.0.2, Ref-router-id 10.255.4.97
  Prefix-count 1
  Prefix feee::10:10:2:0/126
    Prefix-options 0x0, Metric 0
  Aging timer 00:42:20
  Installed 00:17:36 ago, expires in 00:42:21, sent 02:37:54 ago
    OSPF3 AS SCOPE link state database
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Extern      0.0.0.1          10.255.4.85      0x80000002   206  0x9b86  44
  Prefix feee::100:100:1:0/124
  Prefix-options 0x0, Metric 20, Type 2,
  Aging timer 00:56:34
  Installed 00:03:23 ago, expires in 00:56:34, sent 02:37:54 ago
Extern     *0.0.0.1          10.255.4.93      0x80000001  1053  0x59c9  44
  Prefix feee::200:200:1:0/124
  Prefix-options 0x0, Metric 10, Type 2,
  Gen timer 00:25:12
  Aging timer 00:42:26
  Installed 00:17:33 ago, expires in 00:42:27, sent 00:17:31 ago

    OSPF3 Link-Local link state database, interface ge-1/3/0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link       *0.0.0.2          10.255.4.93      0x80000003  1059  0x4dab  64
```

```
  fe80::290:69ff:fe39:1cdb
  Options 0x11, priority 128
  Prefix-count 1
  Prefix feee::10:10:2:0/126 Prefix-options 0x0
  Gen timer 00:12:56
  Aging timer 00:42:20
  Installed 00:17:39 ago, expires in 00:42:21, sent 00:17:37 ago
Link        0.0.0.2          10.255.4.97      0x80000003   205  0xa87d  64
  fe80::290:69ff:fe38:883e
  Options 0x11, priority 128
  Prefix-count 1
  Prefix feee::10:10:2:0/126 Prefix-options 0x0
  Aging timer 00:56:35
  Installed 00:03:22 ago, expires in 00:56:35, sent 02:37:54 ago

    OSPF3 Link-Local link state database, interface so-2/2/0.0
 Type       ID               Adv Rtr           Seq         Age  Cksum  Len
Link        0.0.0.2          10.255.4.85      0x80000002   506  0x42bb  64
  fe80::280:42ff:fe10:f169
  Options 0x13, priority 128
  Prefix-count 1
  Prefix feee::10:10:1:0/126 Prefix-options 0x0
  Aging timer 00:51:34
  Installed 00:08:23 ago, expires in 00:51:34, sent 02:37:54 ago
Link       *0.0.0.3          10.255.4.93      0x80000002   505  0x6b7a  64
  fe80::280:42ff:fe10:f177
  Options 0x13, priority 128
  Prefix-count 1
  Prefix feee::10:10:1:0/126 Prefix-options 0x0
  Gen timer 00:37:28
  Aging timer 00:51:35
  Installed 00:08:25 ago, expires in 00:51:35, sent 00:08:23 ago
  Ours
```

show ospf3 database summary

```
user@host> show ospf3 database summary
Area 0.0.0.0:
   2 Router LSAs
   1 InterArPfx LSAs
   1 InterArRtr LSAs
   1 IntraArPfx LSAs
Area 0.0.0.1:
   2 Router LSAs
   1 Network LSAs
   2 InterArPfx LSAs
   1 NSSA LSAs
   1 IntraArPfx LSAs
Externals:
   2 Extern LSAs
Interface ge-1/3/0.0:
   1 Link LSAs
Interface lo0.0:
Interface so-2/2/0.0:
   1 Link LSAs
```

show ospf database

Syntax

show ospf database
<brief | detail | extensive | summary>
<advertising-router (*address* | self)>
<area *area-id*>
<asbrsummary>
<external>
<instance *instance-name*>
<link-local>
<logical-system (all | *logical-system-name*)>
<lsa-id *lsa-id*>
<netsummary>
<network>
<nssa>
<opaque-area>
<router>

Syntax (J-EX Series Switch)

show ospf database
<brief | detail | extensive | summary>
<advertising-router (*address* | self)>
<area *area-id*>
<asbrsummary>
<external>
<instance *instance-name*>
<link-local>
<lsa-id *lsa-id*>
<netsummary>
<network>
<nssa>
<opaque-area>
<router>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display the entries in the Open Shortest Path First version 2 (OSPFv2) link-state database, which contains data about link-state advertisement (LSA) packets.

Options

none—Display standard information about entries in the OSPFv2 link-state database for all routing instances.

brief | detail | extensive | summary—(Optional) Display the specified level of output.

advertising-router (*address* | self)—(Optional) Display the LSAs advertised either by a particular routing device or by this routing device.

area *area-id*—(Optional) Display the LSAs in a particular area.

asbrsummary—(Optional) Display summary AS boundary router LSA entries.

external—(Optional) Display external LSAs.

instance *instance-name*—(Optional) Display all OSPF database information under the named routing instance.

link-local—(Optional) Display information about link-local LSAs.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

lsa-id *lsa-id*—(Optional) Display the LSA with the specified LSA identifier.

netsummary—(Optional) Display summary network LSAs.

network—(Optional) Display information about network LSAs.

nssa—(Optional) Display information about not-so-stubby area (NSSA) LSAs.

opaque-area—(Optional) Display opaque area-scope LSAs.

router—(Optional) Display information about router LSAs.

Required Privilege Level view

Related Documentation

- **clear (ospf | ospf3) database on page 1748**

List of Sample Output

show ospf database on page 1880
show ospf database brief on page 1881
show ospf database detail on page 1881
show ospf database extensive on page 1882
show ospf database summary on page 1884

Output Fields Table 231 on page 1879 describes the output fields for the **show ospf database** command. Output fields are listed in the approximate order in which they appear.

Table 231: show ospf database Output Fields

Field Name	Field Description	Level of Output
area	Area number. Area **0.0.0.0** is the backbone area.	All levels
Type	Type of link advertisement: **ASBRSum**, **Extern**, **Network**, **NSSA**, **OpaqArea**, **Router**, or **Summary**.	All levels
ID	LSA identifier included in the advertisement. An asterisk preceding the identifier marks database entries that originated from the local routing device.	All levels
Adv Rtr	Address of the routing device that sent the advertisement.	All levels
Seq	Link sequence number of the advertisement.	All levels
Age	Time elapsed since the LSA was originated, in seconds.	All levels
Cksum	Checksum value of the LSA.	All levels

Table 231: show ospf database Output Fields *(continued)*

Field Name	Field Description	Level of Output
Len	Length of the advertisement, in bytes.	All levels
Router	Router link-state advertisement information: • **bits**—Flags describing the routing device that generated the LSP. • **link count**—Number of links in the advertisement. • **id**—ID of a routing device or subnet on the link. • **data**—For stub networks, the subnet mask; otherwise, the IP address of the routing device that generated the LSP. • **type**—Type of link. It can be **PointToPoint**, **Transit**, **Stub**, or **Virtual**. • **TOS count**—Number of type-of-service (ToS) entries in the advertisement. • **TOS 0 metric**—Metric for ToS 0. • **TOS**—Type-of-service (ToS) value. • **metric**—Metric for the ToS.	**detail extensive**
Network	Network link-state advertisement information: • **mask**—Network mask. • **attached router**—ID of the attached neighbor.	**detail extensive**
Summary	Summary link-state advertisement information: • **mask**—Network mask. • **TOS**—Type-of-service (ToS) value. • **metric**—Metric for the ToS.	**detail extensive**
Gen timer	How long until the LSA is regenerated.	**extensive**
Aging time	How long until the LSA expires.	**extensive**
Installed *hh*:*mm*:*ss* ago	How long ago the route was installed.	**extensive**
expires in *hh*:*mm*:*ss*	How long until the route expires.	**extensive**
Ours	Indicates that this is a local advertisement.	**extensive**
Router LSAs	Number of router link-state advertisements in the link-state database.	**summary**
Network LSAs	Number of network link-state advertisements in the link-state database.	**summary**
Summary LSAs	Number of summary link-state advertisements in the link-state database.	**summary**

show ospf database

```
user@host> show ospf database
OSPF link state database, Area 0.0.0.1
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   10.255.70.103    10.255.70.103    0x80000002   215  0x20 0x4112  48
Router  *10.255.71.242    10.255.71.242    0x80000002   214  0x20 0x11b1  48
```

```
Summary *23.1.1.0         10.255.71.242    0x80000002   172  0x20 0x6d72  28
Summary *24.1.1.0         10.255.71.242    0x80000002   177  0x20 0x607e  28
NSSA    *33.1.1.1         10.255.71.242    0x80000002   217  0x28 0x73bd  36

    OSPF link state database, Area 0.0.0.2
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   10.255.71.52     10.255.71.52     0x80000004   174  0x20 0xd021  36
Router  *10.255.71.242    10.255.71.242    0x80000003   173  0x20 0xe191  36
Network *23.1.1.1         10.255.71.242    0x80000002   173  0x20 0x9c76  32
Summary *12.1.1.0         10.255.71.242    0x80000001   217  0x20 0xfeec  28
Summary *24.1.1.0         10.255.71.242    0x80000002   177  0x20 0x607e  28
NSSA    *33.1.1.1         10.255.71.242    0x80000001   222  0x28 0xe047  36

    OSPF link state database, Area 0.0.0.3
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   10.255.71.238    10.255.71.238    0x80000003   179  0x20 0x3942  36
Router  *10.255.71.242    10.255.71.242    0x80000003   177  0x20 0xf37d  36
Network *24.1.1.1         10.255.71.242    0x80000002   177  0x20 0xc591  32
Summary *12.1.1.0         10.255.71.242    0x80000001   217  0x20 0xfeec  28
Summary *23.1.1.0         10.255.71.242    0x80000002   172  0x20 0x6d72  28
NSSA    *33.1.1.1         10.255.71.242    0x80000001   222  0x28 0xeb3b  36
```

show ospf database brief

The output for the **show ospf database brief** command is identical to that for the **show ospf database** command. For sample output, see **show ospf database on page 1880**.

show ospf database detail

```
user@host> show ospf database detail
    OSPF link state database, Area 0.0.0.1
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   10.255.70.103    10.255.70.103    0x80000002   261  0x20 0x4112  48
  bits 0x0, link count 2
  id 10.255.71.242, data 12.1.1.1, Type PointToPoint (1)
  TOS count 0, TOS 0 metric 1
  id 12.1.1.0, data 255.255.255.0, Type Stub (3)
  TOS count 0, TOS 0 metric 1
Router  *10.255.71.242    10.255.71.242    0x80000002   260  0x20 0x11b1  48
  bits 0x3, link count 2
  id 10.255.70.103, data 12.1.1.2, Type PointToPoint (1)
  TOS count 0, TOS 0 metric 1
  id 12.1.1.0, data 255.255.255.0, Type Stub (3)
  TOS count 0, TOS 0 metric 1
Summary *23.1.1.0         10.255.71.242    0x80000002   218  0x20 0x6d72  28
  mask 255.255.255.0
  TOS 0x0, metric 1
Summary *24.1.1.0         10.255.71.242    0x80000002   223  0x20 0x607e  28
  mask 255.255.255.0
  TOS 0x0, metric 1
NSSA    *33.1.1.1         10.255.71.242    0x80000002   263  0x28 0x73bd  36
  mask 255.255.255.255
  Type 2, TOS 0x0, metric 0, fwd addr 12.1.1.2, tag 0.0.0.0

    OSPF link state database, Area 0.0.0.2
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   10.255.71.52     10.255.71.52     0x80000004   220  0x20 0xd021  36
  bits 0x0, link count 1
  id 23.1.1.1, data 23.1.1.2, Type Transit (2)
  TOS count 0, TOS 0 metric 1
Router  *10.255.71.242    10.255.71.242    0x80000003   219  0x20 0xe191  36
  bits 0x3, link count 1
  id 23.1.1.1, data 23.1.1.1, Type Transit (2)
  TOS count 0, TOS 0 metric 1
```

```
Network *23.1.1.1         10.255.71.242    0x80000002   219  0x20 0x9c76  32
  mask 255.255.255.0
  attached router 10.255.71.242
  attached router 10.255.71.52
Summary *12.1.1.0         10.255.71.242    0x80000001   263  0x20 0xfeec  28
  mask 255.255.255.0
  TOS 0x0, metric 1
Summary *24.1.1.0         10.255.71.242    0x80000002   223  0x20 0x607e  28
  mask 255.255.255.0
  TOS 0x0, metric 1
NSSA    *33.1.1.1         10.255.71.242    0x80000001   268  0x28 0xe047  36
  mask 255.255.255.255
  Type 2, TOS 0x0, metric 0, fwd addr 23.1.1.1, tag 0.0.0.0

    OSPF link state database, Area 0.0.0.3
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   10.255.71.238    10.255.71.238    0x80000003   225  0x20 0x3942  36
  bits 0x0, link count 1
  id 24.1.1.1, data 24.1.1.2, Type Transit (2)
  TOS count 0, TOS 0 metric 1
Router  *10.255.71.242    10.255.71.242    0x80000003   223  0x20 0xf37d  36
  bits 0x3, link count 1
  id 24.1.1.1, data 24.1.1.1, Type Transit (2)
  TOS count 0, TOS 0 metric 1
Network *24.1.1.1         10.255.71.242    0x80000002   223  0x20 0xc591  32
  mask 255.255.255.0
  attached router 10.255.71.242
  attached router 10.255.71.238
Summary *12.1.1.0         10.255.71.242    0x80000001   263  0x20 0xfeec  28
  mask 255.255.255.0
  TOS 0x0, metric 1
Summary *23.1.1.0         10.255.71.242    0x80000002   218  0x20 0x6d72  28
  mask 255.255.255.0
  TOS 0x0, metric 1
NSSA    *33.1.1.1         10.255.71.242    0x80000001   268  0x28 0xeb3b  36
  mask 255.255.255.255
  Type 2, TOS 0x0, metric 0, fwd addr 24.1.1.1, tag 0.0.0.0
```

show ospf database extensive

```
user@host> show ospf database extensive
    OSPF link state database, Area 0.0.0.1
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   10.255.70.103    10.255.70.103    0x80000002   286  0x20 0x4112  48
  bits 0x0, link count 2
  id 10.255.71.242, data 12.1.1.1, Type PointToPoint (1)
  TOS count 0, TOS 0 metric 1
  id 12.1.1.0, data 255.255.255.0, Type Stub (3)
  TOS count 0, TOS 0 metric 1
  Aging timer 00:55:14
  Installed 00:04:43 ago, expires in 00:55:14
  Last changed 00:04:43 ago, Change count: 2
Router  *10.255.71.242    10.255.71.242    0x80000002   285  0x20 0x11b1  48
  bits 0x3, link count 2
  id 10.255.70.103, data 12.1.1.2, Type PointToPoint (1)
  TOS count 0, TOS 0 metric 1
  id 12.1.1.0, data 255.255.255.0, Type Stub (3)
  TOS count 0, TOS 0 metric 1
  Gen timer 00:45:15
  Aging timer 00:55:15
  Installed 00:04:45 ago, expires in 00:55:15, sent 00:04:43 ago
  Last changed 00:04:45 ago, Change count: 2, Ours
Summary *23.1.1.0         10.255.71.242    0x80000002   243  0x20 0x6d72  28
```

```
  mask 255.255.255.0
  TOS 0x0, metric 1
  Gen timer 00:45:57
  Aging timer 00:55:57
  Installed 00:04:03 ago, expires in 00:55:57, sent 00:04:01 ago
  Last changed 00:04:48 ago, Change count: 1, Ours
Summary *24.1.1.0         10.255.71.242    0x80000002   248  0x20 0x607e  28
  mask 255.255.255.0
  TOS 0x0, metric 1
  Gen timer 00:45:52
  Aging timer 00:55:52
  Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:06 ago
  Last changed 00:04:48 ago, Change count: 1, Ours
NSSA    *33.1.1.1         10.255.71.242    0x80000002   288  0x28 0x73bd  36
  mask 255.255.255.255
  Type 2, TOS 0x0, metric 0, fwd addr 12.1.1.2, tag 0.0.0.0
  Gen timer 00:45:12
  Aging timer 00:55:12
  Installed 00:04:48 ago, expires in 00:55:12, sent 00:04:48 ago
  Last changed 00:04:48 ago, Change count: 2, Ours

    OSPF link state database, Area 0.0.0.2
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   10.255.71.52     10.255.71.52     0x80000004   245  0x20 0xd021  36
  bits 0x0, link count 1
  id 23.1.1.1, data 23.1.1.2, Type Transit (2)
  TOS count 0, TOS 0 metric 1
  Aging timer 00:55:55
  Installed 00:04:02 ago, expires in 00:55:55
  Last changed 00:04:02 ago, Change count: 2
Router  *10.255.71.242    10.255.71.242    0x80000003   244  0x20 0xe191  36
  bits 0x3, link count 1
  id 23.1.1.1, data 23.1.1.1, Type Transit (2)
  TOS count 0, TOS 0 metric 1
  Gen timer 00:45:56
  Aging timer 00:55:56
  Installed 00:04:04 ago, expires in 00:55:56, sent 00:04:02 ago
  Last changed 00:04:04 ago, Change count: 2, Ours
Network *23.1.1.1         10.255.71.242    0x80000002   244  0x20 0x9c76  32
  mask 255.255.255.0
  attached router 10.255.71.242
  attached router 10.255.71.52
  Gen timer 00:45:56
  Aging timer 00:55:56
  Installed 00:04:04 ago, expires in 00:55:56, sent 00:04:02 ago
  Last changed 00:04:04 ago, Change count: 1, Ours
Summary *12.1.1.0         10.255.71.242    0x80000001   288  0x20 0xfeec  28
  mask 255.255.255.0
  TOS 0x0, metric 1
  Gen timer 00:45:12
  Aging timer 00:55:12
  Installed 00:04:48 ago, expires in 00:55:12, sent 00:04:04 ago
  Last changed 00:04:48 ago, Change count: 1, Ours
Summary *24.1.1.0         10.255.71.242    0x80000002   248  0x20 0x607e  28
  mask 255.255.255.0
  TOS 0x0, metric 1
  Gen timer 00:45:52
  Aging timer 00:55:52
  Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:04 ago
  Last changed 00:04:48 ago, Change count: 1, Ours
NSSA    *33.1.1.1         10.255.71.242    0x80000001   293  0x28 0xe047  36
```

```
  mask 255.255.255.255
  Type 2, TOS 0x0, metric 0, fwd addr 23.1.1.1, tag 0.0.0.0
  Gen timer 00:45:07
  Aging timer 00:55:07
  Installed 00:04:53 ago, expires in 00:55:07, sent 00:04:04 ago
  Last changed 00:04:53 ago, Change count: 1, Ours

    OSPF link state database, Area 0.0.0.3
 Type       ID               Adv Rtr           Seq      Age  Opt  Cksum  Len
Router   10.255.71.238    10.255.71.238    0x80000003   250  0x20 0x3942  36
  bits 0x0, link count 1
  id 24.1.1.1, data 24.1.1.2, Type Transit (2)
  TOS count 0, TOS 0 metric 1
  Aging timer 00:55:50
  Installed 00:04:07 ago, expires in 00:55:50
  Last changed 00:04:07 ago, Change count: 2
Router  *10.255.71.242    10.255.71.242    0x80000003   248  0x20 0xf37d  36
  bits 0x3, link count 1
  id 24.1.1.1, data 24.1.1.1, Type Transit (2)
  TOS count 0, TOS 0 metric 1
  Gen timer 00:45:52
  Aging timer 00:55:52
  Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:06 ago
  Last changed 00:04:08 ago, Change count: 2, Ours
Network *24.1.1.1         10.255.71.242    0x80000002   248  0x20 0xc591  32
  mask 255.255.255.0
  attached router 10.255.71.242
  attached router 10.255.71.238
  Gen timer 00:45:52
  Aging timer 00:55:52
  Installed 00:04:08 ago, expires in 00:55:52, sent 00:04:06 ago
  Last changed 00:04:08 ago, Change count: 1, Ours
Summary *12.1.1.0         10.255.71.242    0x80000001   288  0x20 0xfeec  28
  mask 255.255.255.0
  TOS 0x0, metric 1
  Gen timer 00:45:12
  Aging timer 00:55:12
  Installed 00:04:48 ago, expires in 00:55:12, sent 00:04:13 ago
  Last changed 00:04:48 ago, Change count: 1, Ours
Summary *23.1.1.0         10.255.71.242    0x80000002   243  0x20 0x6d72  28
  mask 255.255.255.0
  TOS 0x0, metric 1
  Gen timer 00:45:57
  Aging timer 00:55:57
  Installed 00:04:03 ago, expires in 00:55:57, sent 00:04:01 ago
  Last changed 00:04:48 ago, Change count: 1, Ours
NSSA    *33.1.1.1         10.255.71.242    0x80000001   293  0x28 0xeb3b  36
  mask 255.255.255.255
  Type 2, TOS 0x0, metric 0, fwd addr 24.1.1.1, tag 0.0.0.0
  Gen timer 00:45:07
  Aging timer 00:55:07
  Installed 00:04:53 ago, expires in 00:55:07, sent 00:04:13 ago
  Last changed 00:04:53 ago, Change count: 1, Ours
```

show ospf database summary

```
user@host> show ospf database summary
Area 0.0.0.1:
   2 Router LSAs
   2 Summary LSAs
   1 NSSA LSAs
Area 0.0.0.2:
   2 Router LSAs
```

```
   1 Network LSAs
   2 Summary LSAs
   1 NSSA LSAs
Area 0.0.0.3:
   2 Router LSAs
   1 Network LSAs
   2 Summary LSAs
   1 NSSA LSAs
Externals:
Interface fe-2/2/1.0:
Interface ge-0/3/2.0:
Interface so-0/1/2.0:
Interface so-0/1/2.0:
```

show policy damping

Syntax
show policy damping
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show policy damping

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display information about Border Gateway Protocol (BGP) route flap damping parameters.

Options
none—Display information about BGP route flap damping parameters.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information
In the output from this command, figure-of-merit values correlate to the probability of future instability of a routing device. Routes with higher figure-of-merit values are suppressed for longer periods of time. The figure-of-merit value decays exponentially over time. A figure-of-merit value of zero is assigned to each new route. The value is increased each time the route is withdrawn or readvertised, or when one of its path attributes changes.

Required Privilege Level
view

Related Documentation
- "Configuring BGP Flap Damping Parameters" in the *Junos OS Policy Framework Configuration Guide*
- **clear bgp damping on page 1755**
- **show route damping on page 1922**

List of Sample Output
show policy damping on page 1887

Output Fields
Table 232 on page 1886 describes the output fields for the **show policy damping** command. Output fields are listed in the approximate order in which they appear.

Table 232: show policy damping Output Fields

Field Name	Field Description
Halflife	Decay half-life, in minutes. The value represents the period during which the accumulated figure-of-merit value is reduced by half if the route remains stable. If a route has flapped, but then becomes stable, the figure-of-merit value for the route decays exponentially. For example, for a route with a figure-of-merit value of 1500, if no incidents occur, its figure-of-merit value is reduced to 750 after 15 minutes and to 375 after another 15 minutes.

Table 232: show policy damping Output Fields *(continued)*

Field Name	Field Description
Reuse merit	Figure-of-merit value below which a suppressed route can be used again. A suppressed route becomes reusable when its figure-of-merit value decays to a value below a reuse threshold, and the route once again is considered usable and can be installed in the forwarding table and exported from the routing table.
Suppress/cutoff merit	Figure-of-merit value above which a route is suppressed for use or inclusion in advertisements. When a route's figure-of-merit value reaches a particular level, called the cutoff or suppression threshold, the route is suppressed. When a route is suppressed, the routing table no longer installs the route into the forwarding table and no longer exports this route to any of the routing protocols.
Maximum suppress time	Maximum hold-down time, in minutes. The value represents the maximum time that a route can be suppressed no matter how unstable it has been before this period of stability.
Computed values	• **Merit ceiling**—Maximum merit that a flapping route can collect. • **Maximum decay**—Maximum decay half-life, in minutes.

show policy damping

```
user@host> show policy damping
Default damping information:
  Halflife: 15 minutes
  Reuse merit: 750 Suppress/cutoff merit: 3000
  Maximum suppress time: 60 minutes
  Computed values:
    Merit ceiling: 12110
    Maximum decay: 6193
Damping information for "standard-damping":
  Halflife: 10 minutes
  Reuse merit: 4000 Suppress/cutoff merit: 8000
  Maximum suppress time: 30 minutes
  Computed values:
    Merit ceiling: 32120
    Maximum decay: 12453
```

show rip general-statistics

Syntax
show rip general-statistics
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show rip general-statistics

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display brief Routing Information Protocol (RIP) statistics.

Options
none—Display brief RIP statistics.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- **clear rip general-statistics on page 1769**

List of Sample Output
show rip general-statistics on page 1888

Output Fields
Table 233 on page 1888 lists the output fields for the **show rip general-statistics** command. Output fields are listed in the approximate order in which they appear.

Table 233: show rip general-statistics Output Fields

Field Name	Field Description
bad msgs	Number of invalid messages received.
no recv intf	Number of packets received with no matching interface.
curr memory	Amount of memory currently used by RIP.
max memory	Most memory used by RIP.

show rip general-statistics

```
user@host> show rip general-statistics
RIPv2 I/O info:
     bad msgs       :         0
     no recv intf   :         0
     curr memory    :         0
     max memory     :         0
```

show rip neighbor

Syntax

show rip neighbor
<instance (all | *instance-name*)>
<logical-system (all | *logical-system-name*)>
<*name*>

Syntax (J-EX Series Switch)

show rip neighbor
<instance (all | *instance-name*)>
<*name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display information about Routing Information Protocol (RIP) neighbors.

Options

none—Display information about all RIP neighbors for all instances.

instance (all | *instance-name*)—(Optional) Display RIP neighbor information for all instances or for only the specified routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

name—(Optional) Display detailed information about only the specified RIP neighbor.

Required Privilege Level

view

List of Sample Output

show rip neighbor on page 1890

Output Fields

Table 234 on page 1889 lists the output fields for the **show rip neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 234: show rip neighbor Output Fields

Field Name	Field Description
Neighbor	Name of RIP neighbor.
State	State of the connection: **Up** or **Dn** (Down).
Source Address	Source address.
Destination Address	Destination address.
Send Mode	Send options: **broadcast**, **multicast**, **none**, or **version 1**.
Receive Mode	Type of packets to accept: **both**, **none**, **version 1**, or **version 2**.
In Met	Metric added to incoming routes when advertising into RIP routes that were learned from other protocols.

show rip neighbor `user@host>` **show rip neighbor**

```
                                Source          Destination     Send    Receive  In
Neighbor            State       Address         Address         Mode    Mode     Met
--------            -----       -------         -----------     ----    -------  ---
ge-2/3/0.0             Up       192.168.9.105   192.168.9.107   bcast   both       1
at-5/1/1.42            Dn       (null)          (null)          mcast   v2 only    3
at-5/1/0.42            Dn       (null)          (null)          mcast   both       3
at-5/1/0.0             Up       20.0.0.1        224.0.0.9       mcast   both       3
so-0/0/0.0             Up       192.168.9.97    224.0.0.9       mcast   both       3
```

show rip statistics

Syntax
show rip statistics
<instance (all | *instance-name*)>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show rip statistics
<instance (all | *instance-name*)>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display Routing Information Protocol (RIP) statistics about messages sent and received on an interface, as well as information received from advertisements from other routing devices.

Options
none—Display RIP statistics for all routing instances.

instance (all | *instance-name*—(Optional) Display RIP statistics for all instances or for only the specified routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- **clear rip statistics on page 1770**

List of Sample Output
show rip statistics on page 1892

Output Fields
Table 235 on page 1891 lists the output fields for the **show rip statistics** command. Output fields are listed in the approximate order in which they appear.

Table 235: show rip statistics Output Fields

Field Name	Field Description
RIP info	Information about RIP on the specified interface: • **port**—UDP port number used for RIP. • **holddown**—Hold-down interval, in seconds. • **rts learned**—Number of routes learned through RIP. • **rts held down**—Number of routes held down by RIP. • **rqsts dropped**—Number of received request packets that were dropped. • **resps dropped**—Number of received response packets that were dropped. • **restart**—Graceful restart status. Displayed when RIP is or has been in the process of graceful restart.

Table 235: show rip statistics Output Fields *(continued)*

Field Name	Field Description
logical-interface	Name of the logical interface and its statistics: • **routes learned**—Number of routes learned on the logical interface. • **routes advertised**—Number of routes advertised by the logical interface. • **timeout**—Timeout interval, in seconds. • **update interval**—Number of seconds since last update.
Counter	List of counter types: • **Updates Sent**—Number of update messages sent. • **Triggered Updates Sent**—Number of triggered update messages sent. • **Responses Sent**—Number of response messages sent. • **Bad Messages**—Number of invalid messages received. • **RIPv1 Updates Received**—Number of RIPv1 update messages received. • **RIPv1 Bad Route Entries**—Number of RIPv1 invalid route entry messages received. • **RIPv1 Updates Ignored**—Number of RIPv1 update messages ignored. • **RIPv2 Updates Received**—Number of RIPv2 update messages received. • **RIPv2 Bad Route Entries**—Number of RIPv2 invalid route entry messages received. • **RIPv2 Updates Ignored**—Number of RIPv2 update messages that were ignored. • **Authentication Failures**—Number of received update messages that failed authentication. • **RIP Requests Received**—Number of RIP request messages received. • **RIP Requests Ignored**—Number of RIP request messages ignored.
Total	Total number of packets for the selected counter.
Last 5 min	Number of packets for the selected counter in the most recent 5-minute period.
Last minute	Number of packets for the selected counter in the most recent 1-minute period.

show rip statistics

```
user@host> show rip statistics so-0/0/0.0
RIP info: port 520; update interval: 30s; holddown 180s; timeout 120s
restart in progress: restart time 60s; restart will complete in 55s
    rts learned  rts held down  rqsts dropped  resps dropped
              0              0               0              0
so-0/0/0.0:  0 routes learned; 501 routes advertised
Counter                         Total   Last 5 min  Last minute
-------                   -----------  -----------  -----------
Updates Sent                        0            0            0
Triggered Updates Sent              0            0            0
Responses Sent                      0            0            0
Bad Messages                        0            0            0
RIPv1 Updates Received              0            0            0
RIPv1 Bad Route Entries             0            0            0
RIPv1 Updates Ignored               0            0            0
RIPv2 Updates Received              0            0            0
RIPv2 Bad Route Entries             0            0            0
RIPv2 Updates Ignored               0            0            0
Authentication Failures             0            0            0
```

```
RIP Requests Received              0            0            0
RIP Requests Ignored               0            0            0
```

show ripng general-statistics

Syntax
show ripng general-statistics
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show ripng general-statistics

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display general Routing Information Protocol next-generation (RIPng) statistics.

Options
none—Display general RIPng statistics.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- **clear ripng general-statistics on page 1771**

List of Sample Output
show ripng general-statistics on page 1894

Output Fields
Table 236 on page 1894 lists the output fields for the **show ripng general-statistics** command. Output fields are listed in the approximate order in which they appear.

Table 236: show ripng general-statistics Output Fields

Field Name	Field Description
bad msgs	Number of invalid messages received.
no recv intf	Number of packets received with no matching interface.
curr memory	Amount of memory currently used by RIPng.
max memory	Most memory used by RIPng.

show ripng general-statistics

```
user@host> show ripng general-statistics
RIPng I/O info:
     bad msgs       :         0
     no recv intf   :         0
     curr memory    :         0
     max memory     :         0
```

show ripng neighbor

Syntax

show ripng neighbor
<logical-system (all | *logical-system-name*)>
<*name*>

Syntax (J-EX Series Switch)

show ripng neighbor
<*name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display information about Routing Information Protocol next-generation (RIPng) neighbors.

Options

none—Display information about all RIPng neighbors.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

name—(Optional) Display detailed information about a specific RIPng neighbor.

Required Privilege Level

view

List of Sample Output

show ripng neighbor on page 1895

Output Fields

Table 237 on page 1895 lists the output fields for the **show ripng neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 237: show ripng neighbor Output Fields

Field Name	Field Description
Neighbor	Name of RIPng neighbor.
State	State of the connection: **Up** or **Dn** (Down).
Source Address	Source address.
Destination Address	Destination address.
Send Mode	Send options: **broadcast**, **multicast**, **none**, **version 1**, or **yes.**
Receive Mode	Type of packets to accept: **both**, **none**, **version 1**, or **yes.**
In Met	Metric added to incoming routes when advertising into RIPng routes that were learned from other protocols.

show ripng neighbor

```
user@host> show ripng neighbor
                     Source                                  Dest              In
Neighbor     State   Address                                 Address  Send Recv Met
```

```
--------     -----  -------                          -------  ---- ---- ---
fe-0/0/2.0      Up  fe80::290:69ff:fe68:b002         ff02::9  yes  yes    1
```

show ripng statistics

Syntax
show ripng statistics
<logical-system (all | *logical-system-name*)>
<name>

Syntax (J-EX Series Switch)
show ripng statistics
<name>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display Routing Information Protocol next generation (RIPng) statistics about messages sent and received on an interface, as well as information received from advertisements from other routing devices.

Options
none—Display RIPng statistics for all neighbors.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

name—(Optional) Display detailed information about a specific RIPng neighbor.

Required Privilege Level
view

Related Documentation
- **clear ripng statistics on page 1772**

List of Sample Output
show ripng statistics on page 1898

Output Fields
Table 238 on page 1897 lists the output fields for the **show ripng statistics** command. Output fields are listed in the approximate order in which they appear.

Table 238: show ripng statistics Output Fields

Field Name	Field Description
RIPng info	Information about RIPng on the specified interface: • **port**—UDP port number used for RIP. • **holddown**—Hold-down interval, in seconds. • **rts learned**—Number of routes learned through RIP. • **rts held down**—Number of routes held down by RIP. • **rqsts dropped**—Number of received request packets that were dropped. • **resps dropped**—Number of received response packets that were dropped. • **restart**—Graceful restart status. Displayed when RIPng is or has been in the process of graceful restart.

Table 238: show ripng statistics Output Fields *(continued)*

Field Name	Field Description
logical-interface	Name of the logical interface and its statistics: • **routes learned**—Number of routes learned on the logical interface. • **routes advertised**—Number of routes advertised by the logical interface. • **timeout**—Timeout interval, in seconds. • **update interval**—Number of seconds since last update.
Counter	List of counter types: • **Updates Sent**—Number of update messages sent. • **Triggered Updates Sent**—Number of triggered update messages sent. • **Responses Sent**—Number of response messages sent. • **Bad Messages**—Number of invalid messages received. • **Updates Received**—Number of RIPng update messages received. • **Bad Route Entries**—Number of RIPng invalid route entry messages received. • **Updates Ignored**—Number of RIPng update messages ignored. • **RIPng Requests Received**—Number of RIPng request messages received. • **RIPng Requests Ignored**—Number of RIPng request messages ignored.
Total	Total number of packets for the selected counter.
Last 5 min	Number of packets for the selected counter in the most recent 5-minute period.
Last minute	Number of packets for the selected counter in the most recent 1-minute period.

show ripng statistics

```
user@host> show ripng statistics
RIPng info: port 521; holddown 120s;
    rts learned  rts held down  rqsts dropped  resps dropped
              0              0              0              0

so-0/1/3.0:  0 routes learned; 1 routes advertised; timeout 180s; update interval
 20s
Counter                       Total   Last 5 min  Last minute
-------                   -----------  ----------  -----------
Updates Sent                      934          16            4
Triggered Updates Sent              1           0            0
Responses Sent                      0           0            0
Bad Messages                        0           0            0
Updates Received                    0           0            0
Bad Route Entries                   0           0            0
Updates Ignored                     0           0            0
RIPng Requests Received             0           0            0
RIPng Requests Ignored              0           0            0
```

show route

Syntax

show route
<all>
<*destination-prefix*>
<logical-system (all | *logical-system-name*)>
<private>

Syntax (J-EX Series Switch)

show route
<all>
<*destination-prefix*>
<private>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display the active entries in the routing tables.

Options

none—Display brief information about all active entries in the routing tables.

all—(Optional) Display information about all routing tables, including private, or internal, routing tables.

destination-prefix—(Optional) Display active entries for the specified address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

private—(Optional) Display information only about all private, or internal, routing tables.

Required Privilege Level

view

List of Sample Output

show route on page 1902
show route destination-prefix on page 1902

Output Fields

Table 239 on page 1899 describes the output fields for the **show route** command. Output fields are listed in the approximate order in which they appear.

Table 239: show route Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, **inet.0**).
number **destinations**	Number of destinations for which there are routes in the routing table.
number **routes**	Number of routes in the routing table and total number of routes in the following states: • **active** (routes that are active). • **holddown** (routes that are in the pending state before being declared inactive). • **hidden** (routes that are not used because of a routing policy).

Table 239: show route Output Fields *(continued)*

Field Name	Field Description

destination-prefix	Route destination (for example:**10.0.0.1/24**). Sometimes the route information is presented in another format, such as:
	• ***MPLS-label*** (for example, **80001**).
	• ***interface-name*** (for example, **ge-1/0/2**).
	• ***neighbor-address*:*control-word-status*:*encapsulation type*:*vc-id* :*source*** (Layer 2 circuit only; for example, **10.1.1.195:NoCtrlWord:1:1:Local/96**):
	  ◦ ***neighbor-address***—Address of the neighbor.
	  ◦ ***control-word-status***—Whether the use of the control word has been negotiated for this virtual circuit: **NoCtrlWord** or **CtrlWord**.
	  ◦ ***encapsulation type***—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.
	  ◦ ***vc-id***—Virtual circuit identifier.
	• ***source***—Source of the advertisement: **Local** or **Remote**.

[***protocol*, *preference*** **]**	Protocol from which the route was learned and the preference value for the route. • **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • **-** —A hyphen indicates the last active route. • *****—An asterisk indicates that the route is both the active and the last active route. An asterisk before a **to** line indicates the best subpath to the route. In every routing metric except for the BGP **LocalPref** attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the **LocalPref** value in the **Preference2** field. For example, if the **LocalPref** value for Route 1 is 100, the **Preference2** value is -101. If the **LocalPref** value for Route 2 is 155, the **Preference2** value is -156. Route 2 is preferred because it has a higher **LocalPref** value and a lower **Preference2** value.
weeks:days hours*:*minutes*:*seconds	How long the route been known (for example, **2w4d 13:11:14**, or 2 weeks, 4 days, 13 hours, 11 minutes and 14 seconds).
metric	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
localpref	Local preference value included in the route.
from	Interface from which the route was received.

Table 239: show route Output Fields *(continued)*

Field Name	Field Description
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set.
to	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word **Selected**. This field can also contain the following information: • **Weight**—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when Multiprotocol Label Switching (MPLS) label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • **Balance**—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable Border Gateway Protocol (BGP) multipath load balancing. • ***lsp-path-name***—Name of the label-switched path (LSP) used to reach the next hop. • ***label-action***—MPLS label and operation occurring at the next hop. The operation can be **pop** (where a label is removed from the top of the stack), **push** (where another label is added to the label stack), or **swap** (where a label is replaced by another label).

show route

```
user@host> show route
inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
0.0.0.0/0          *[Static/5] 1w5d 20:30:29
                      Discard
10.255.245.51/32   *[Direct/0] 2w4d 13:11:14
                    > via lo0.0
172.16.0.0/12      *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0
192.168.0.0/18     *[Static/5] 1w5d 20:30:29
                    > to 192.168.167.254 via fxp0.0
192.168.40.0/22    *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0
192.168.64.0/18    *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0
192.168.164.0/22   *[Direct/0] 2w4d 13:11:14
                    > via fxp0.0
192.168.164.51/32  *[Local/0] 2w4d 13:11:14
                      Local via fxp0.0
207.17.136.192/32  *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0

green.inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
100.101.0.0/16     *[Direct/0] 1w5d 20:30:28
                    > via fe-0/0/3.0
100.101.2.3/32     *[Local/0] 1w5d 20:30:28
                      Local via fe-0/0/3.0
224.0.0.5/32       *[OSPF/10] 1w5d 20:30:29, metric 1
                      MultiRecv

red.inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.10.10.10/32     *[Direct/0] 01:08:46
                    > via lo0.1
10.255.245.212/32  *[BGP/170] 00:01:40, localpref 100, from 10.255.245.204
                      AS path: 300 I
                    > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
10.255.245.213/32  *[BGP/170] 00:40:47, localpref 100
                      AS path: 100 I
                    > to 100.1.1.1 via so-0/0/1.0
```

show route destination-prefix

```
user@host> show route 172.16.0.0/12

inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.0.0/12      *[Static/5] 2w4d 12:54:27
                    > to 192.168.167.254 via fxp0.0
```

show route active-path

Syntax
show route active-path
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show route active-path
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display all active routes for destinations. An active route is a route that is selected as the best path. Inactive routes are not displayed.

Options
none—Display all active routes.

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route active-path on page 1903
show route active-path brief on page 1904
show route active-path detail on page 1904
show route active-path extensive on page 1905
show route active-path terse on page 1906

Output Fields
For information about output fields, see the output field tables for the **show route** command, the **show route detail**, the **show route extensive**, or the **show route terse**.

show route active-path

```
user@host> show route active-path

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.70.19/32    *[Direct/0] 21:33:52
                    > via lo0.0
10.255.71.50/32    *[IS-IS/15] 00:18:13, metric 10
                    > to 100.1.2.1 via so-2/1/3.0
100.1.2.0/24       *[Direct/0] 00:18:36
                    > via so-2/1/3.0
100.1.2.2/32       *[Local/0] 00:18:41
                      Local via so-2/1/3.0
192.168.64.0/21    *[Direct/0] 21:33:52
                    > via fxp0.0
192.168.70.19/32   *[Local/0] 21:33:52
                      Local via fxp0.0
```

show route active-path brief

The output for the **show route active-path brief** command is identical to that for the **show route active-path** command. For sample output, see **show route active-path on page 1903**.

show route active-path detail

```
user@host> show route active-path detail

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)

10.255.70.19/32 (1 entry, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via lo0.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 21:37:10
                Task: IF
                Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

                AS path: I

10.255.71.50/32 (1 entry, 1 announced)
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Router, Next hop index: 397
                Next-hop reference count: 4
                Next hop: 100.1.2.1 via so-2/1/3.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 21:31      Metric: 10
                Task: IS-IS
                Announcement bits (4): 0-KRT 2-IS-IS 5-Resolve tree 2 6-Resolve
tree 3
                AS path: I

100.1.2.0/24 (1 entry, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via so-2/1/3.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 21:54
                Task: IF
                Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

                AS path: I

100.1.2.2/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 11
                Interface: so-2/1/3.0
                State: ‹Active NoReadvrt Int›
                Local AS:   200
                Age: 21:59
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I

192.168.64.0/21 (1 entry, 1 announced)
```

```
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via fxp0.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 21:37:10
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I

192.168.70.19/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 11
                Interface: fxp0.0
                State: ‹Active NoReadvrt Int›
                Local AS:   200
                Age: 21:37:10
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I
```

show route active-path extensive

```
user@host> show route active-path extensive

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
10.255.70.19/32 (1 entry, 1 announced)
TSI:
IS-IS level 1, LSP fragment 0
IS-IS level 2, LSP fragment 0
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via lo0.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 21:39:47
                Task: IF
                Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

                AS path: I

10.255.71.50/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.255.71.50/32 -> {100.1.2.1}
IS-IS level 2, LSP fragment 0
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Router, Next hop index: 397
                Next-hop reference count: 4
                Next hop: 100.1.2.1 via so-2/1/3.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 24:08      Metric: 10
                Task: IS-IS
                Announcement bits (4): 0-KRT 2-IS-IS 5-Resolve tree 2 6-Resolve
tree 3
                AS path: I

100.1.2.0/24 (1 entry, 1 announced)
TSI:
```

```
IS-IS level 1, LSP fragment 0
IS-IS level 2, LSP fragment 0
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via so-2/1/3.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 24:31
                Task: IF
                Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

                AS path: I

100.1.2.2/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 11
                Interface: so-2/1/3.0
                State: ‹Active NoReadvrt Int›
                Local AS:   200
                Age: 24:36
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I

192.168.64.0/21 (1 entry, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via fxp0.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 21:39:47
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I

192.168.70.19/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 11
                Interface: fxp0.0
                State: ‹Active NoReadvrt Int›
                Local AS:   200
                Age: 21:39:47
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I
```

show route active-path terse

```
user@host> show route active-path terse

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.255.70.19/32    D   0                     >lo0.0
* 10.255.71.50/32    I  15         10          >100.1.2.1
* 100.1.2.0/24       D   0                     >so-2/1/3.0
* 100.1.2.2/32       L   0                      Local
```

```
* 192.168.64.0/21    D   0                    >fxp0.0
* 192.168.70.19/32   L   0                     Local
```

show route all

Syntax
show route all
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show route all

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display information about all routes in all routing tables, including private, or internal, tables.

Options
none—Display information about all routes in all routing tables, including private, or internal, tables.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route all on page 1908

Output Fields
The output fields for the **show route all** command display all routing tables, including private, or hidden, routing tables. The output field table of the **show route** command does not display entries for private, or hidden, routing tables.

show route all
The following example displays a snippet of output from the **show route** command and then displays the same snippet of output from the **show route all** command:

```
user@host> show route
mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
0                  *[MPLS/0] 2d 02:24:39, metric 1
                      Receive
1                  *[MPLS/0] 2d 02:24:39, metric 1
                      Receive
2                  *[MPLS/0] 2d 02:24:39, metric 1
                      Receive
800017             *[VPLS/7] 1d 14:00:16
                    > via vt-3/2/0.32769, Pop
800018             *[VPLS/7] 1d 14:00:26
                    > via vt-3/2/0.32772, Pop

user@host> show route all
mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
0                  *[MPLS/0] 2d 02:19:12, metric 1
                      Receive
1                  *[MPLS/0] 2d 02:19:12, metric 1
                      Receive
2                  *[MPLS/0] 2d 02:19:12, metric 1
```

```
                    Receive
800017             *[VPLS/7] 1d 13:54:49
                     > via vt-3/2/0.32769, Pop
800018             *[VPLS/7] 1d 13:54:59
                     > via vt-3/2/0.32772, Pop
vt-3/2/0.32769      [VPLS/7] 1d 13:54:49
                      Unusable
vt-3/2/0.32772      [VPLS/7] 1d 13:54:59
                      Unusable
```

show route aspath-regex

Syntax

show route aspath-regex *regular-expression*
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)

show route aspath-regex *regular-expression*

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display the entries in the routing table that match the specified autonomous system (AS) path regular expression.

Options

regular-expression—Regular expression that matches an entire AS path.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information

You can specify a regular expression as:

- An individual AS number
- A period wildcard used in place of an AS number
- An AS path regular expression that is enclosed in parentheses

You also can include the operators described in the table of AS path regular expression operators in the *Junos OS Policy Framework Configuration Guide*. The following list summarizes these operators:

- ***{m,n}***—At least *m* and at most *n* repetitions of the AS path term.
- ***{m}***—Exactly *m* repetitions of the AS path term.
- ***{m,}***—*m* or more repetitions of the AS path term.
- *******—Zero or more repetitions of an AS path term.
- ***+***—One or more repetitions of an AS path term.
- ***?***—Zero or one repetition of an AS path term.
- ***aspath_term | aspath_term***—Match one of the two AS path terms.

  When you specify more than one AS number or path term, or when you include an operator in the regular expression, enclose the entire regular expression in quotation marks. For example, to match any path that contains AS number 234, specify the following command:

  ```
  show route aspath-regex ".* 234 .*"
  ```

Required Privilege Level

view

List of Sample Output

Output Fields

For information about output fields, see the output field table for the **show route** command.

show route aspath-regex (Matching a Specific AS Number)

```
user@host> show route aspath-regex 65477
inet.0: 46411 destinations, 46411 routes (46409 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

111.222.1.0/25     *[BGP/170] 00:08:48, localpref 100, from 111.222.2.24
                   AS Path: [65477] ({65488 65535}) IGP
                     to 111.222.18.225 via fpa0.0(111.222.18.233)
111.222.1.128/25   *[IS-IS/15] 09:15:37, metric 37, tag 1
                     to 111.222.18.225 via fpa0.0(111.222.18.233)
                    [BGP/170] 00:08:48, localpref 100, from 111.222.2.24
                   AS Path: [65477] ({65488 65535}) IGP
                     to 111.222.18.225 via fpa0.0(111.222.18.233)
...
```

show route aspath-regex (Matching Any Path with Two AS Numbers)

```
user@host> show route aspath-regex ?.* 234 3561 .*?

inet.0: 46351 destinations, 46351 routes (46349 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

9.20.0.0/17        *[BGP/170] 01:35:00, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 2685 2686 Incomplete
                     to 192.156.169.1 via 192.156.169.14(so-0/0/0)
12.10.231.0/24     *[BGP/170] 01:35:00, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 5696 7369 IGP
                     to 192.156.169.1 via 192.156.169.14(so-0/0/0)
24.64.32.0/19      *[BGP/170] 01:34:59, localpref 100, from 131.103.20.49
                   AS Path: [666] 234 3561 6327 IGP
                     to 192.156.169.1 via 192.156.169.14(so-0/0/0)
...
```

show route best

Syntax
show route best *destination-prefix*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show route best *destination-prefix*
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display the route in the routing table that is the best route to the specified address or range of addresses. The best route is the longest matching route.

Options
destination-prefix—Address or range of addresses.

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route best on page 1912
show route best detail on page 1914
show route best extensive on page 1914
show route best terse on page 1915

Output Fields
For information about output field tables for the **show route** command, the **show route detail** command, the **show route extensive** command, or the **show route terse** command.

show route best

```
user@host> show route best 10.255.70.103
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.255.70.103/32   *[OSPF/10] 1d 13:19:20, metric 2
                    > to 10.31.1.6 via ge-3/1/0.0
                      via so-0/3/0.0

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.255.70.103/32   *[RSVP/7] 1d 13:20:13, metric 2
                    > via so-0/3/0.0, label-switched-path green-r1-r3

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.0.0.0/8         *[Direct/0] 2d 01:43:34
                    > via fxp2.0
```

```
[Direct/0] 2d 01:43:34
> via fxp1.0
```

show route best detail

```
user@host> show route best 10.255.70.103 detail
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
10.255.70.103/32 (1 entry, 1 announced)
        *OSPF   Preference: 10
                Next-hop reference count: 9
                Next hop: 10.31.1.6 via ge-3/1/0.0, selected
                Next hop: via so-0/3/0.0
                State: <Active Int>
                Local AS:    69
                Age: 1d 13:20:06        Metric: 2
                Area: 0.0.0.0
                Task: OSPF
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
10.255.70.103/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 5
                Next hop: via so-0/3/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 100016
                State: <Active Int>
                Local AS:    69
                Age: 1d 13:20:59        Metric: 2
                Task: RSVP
                Announcement bits (1): 1-Resolve tree 2
                AS path: I

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
10.0.0.0/8 (2 entries, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via fxp2.0, selected
                State: <Active Int>
                Age: 2d 1:44:20
                Task: IF
                AS path: I
         Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via fxp1.0, selected
                State: <NotBest Int>
                Inactive reason: No difference
                Age: 2d 1:44:20
                Task: IF
                AS path: I
```

show route best extensive

The output for the **show route best extensive** command is identical to that for the **show route best detail** command. For sample output, see the **show route best detail on page 1914**.

show route best terse

```
user@host> show route best 10.255.70.103 terse
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.255.70.103/32   O  10          2            >10.31.1.6
                                                   so-0/3/0.0

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.255.70.103/32   R   7          2            >so-0/3/0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.0.0.0/8         D   0                      >fxp2.0
                     D   0                      >fxp1.0
```

show route brief

Syntax
show route brief
<*destination-prefix*>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show route brief
<*destination-prefix*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display brief information about the active entries in the routing tables.

Options
none—Display all active entries in the routing table.

destination-prefix—(Optional) Display active entries for the specified address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route brief on page 1916

Output Fields
For information about output fields, see the Output Field table of the **show route** command.

show route brief

```
user@host> show route brief
inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 1w5d 20:30:29
                      Discard
10.255.245.51/32   *[Direct/0] 2w4d 13:11:14
                    > via lo0.0
172.16.0.0/12      *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0
192.168.0.0/18     *[Static/5] 1w5d 20:30:29
                    > to 192.168.167.254 via fxp0.0
192.168.40.0/22    *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0
192.168.64.0/18    *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0
192.168.164.0/22   *[Direct/0] 2w4d 13:11:14
                    > via fxp0.0
192.168.164.51/32  *[Local/0] 2w4d 13:11:14
                      Local via fxp0.0
207.17.136.192/32  *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0
```

```
green.inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
100.101.0.0/16     *[Direct/0] 1w5d 20:30:28
                    > via fe-0/0/3.0
100.101.2.3/32     *[Local/0] 1w5d 20:30:28
                      Local via fe-0/0/3.0
224.0.0.5/32       *[OSPF/10] 1w5d 20:30:29, metric 1
                      MultiRecv
```

show route community

Syntax
show route community *as-number:community-value*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show route community *as-number:community-value*
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display the route entries in each routing table that are members of a Border Gateway Protocol (BGP) community.

Options
as-number:community-value—One or more community identifiers. ***as-number*** is the AS number, and ***community-value*** is the community identifier. When you specify more than one community identifier, enclose the identifiers in double quotation marks. Community identifiers can include wildcards.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information
Specifying the community option displays all routes matching the community found within the routing table. The community option does not limit the output to only the routes being advertised to the neighbor after any egress routing policy.

Required Privilege Level
view

Related Documentation
- **show route detail on page 1927**

List of Sample Output
show route community on page 1918

Output Fields
For information about output fields, see the output field tables for the **show route** command, the **show route detail** command, the **show route extensive** command, or the **show route terse** command.

show route community

```
user@host> show route community 234:80
inet.0: 46511 destinations, 46511 routes (46509 active, 0 holddown, 2 hidden)
+ = Active Route, - = Last Active, * = Both

4.0.0.0/8          *[BGP/170] 03:33:07, localpref 100, from 131.103.20.49
                   AS Path: {666} 234 2548 1 IGP
                     to 192.156.169.1 via 192.156.169.14(so-0/0/0)
6.0.0.0/8          *[BGP/170] 03:33:07, localpref 100, from 131.103.20.49
                   AS Path: {666} 234 2548 568 721 Incomplete
                     to 192.156.169.1 via 192.156.169.14(so-0/0/0)
9.2.0.0/16         *[BGP/170] 03:33:06, localpref 100, from 131.103.20.49
```

```
AS Path: {666} 234 2548 1673 1675 1747 IGP
  to 192.156.169.1 via 192.156.169.14(so-0/0/0)
```

show route community-name

Syntax

show route community-name *community-name*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)

show route community-name *community-name*
<brief | detail | extensive | terse>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display the route entries in each routing table that are members of a Border Gateway Protocol (BGP) community, specified by a community name.

Options

community-name—Name of the community.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show route community-name on page 1920

Output Fields

For information about output fields, see the output field tables for the **show route** command, the **show route detail** command, the **show route extensive** command, or the **show route terse** command.

show route community-name

```
user@host> show route community-name red-com
inet.0: 17 destinations, 17 routes (16 active, 0 holddown, 1 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

instance1.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.245.212/32  *[BGP/170] 00:04:40, localpref 100, from 10.255.245.204
                      AS path: 300 I
                    > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
20.20.20.20/32     *[BGP/170] 00:04:40, localpref 100, from 10.255.245.204
                      AS path: I
                    > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
100.1.4.0/24       *[BGP/170] 00:04:40, localpref 100, from 10.255.245.204
                      AS path: I
                    > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

bgp.l3vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
```

```
+ = Active Route, - = Last Active, * = Both

10.255.245.204:10:10.255.245.212/32
                   *[BGP/170] 00:06:40, localpref 100, from 10.255.245.204
                      AS path: 300 I
                    > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
10.255.245.204:10:20.20.20.20/32
                   *[BGP/170] 00:36:02, localpref 100, from 10.255.245.204
                      AS path: I
                    > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix
10.255.245.204:10:100.1.4.0/24
                   *[BGP/170] 00:36:02, localpref 100, from 10.255.245.204
                      AS path: I
                    > to 100.1.2.2 via ge-1/1/0.0, label-switched-path to_fix

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

instance1.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route damping

Syntax
show route damping (decayed | history | suppressed)
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show route damping (decayed | history | suppressed)
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display the Border Gateway Protocol (BGP) routes for which updates might have been reduced because of route flap damping.

Options
brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.

decayed—Display route damping entries that might no longer be valid, but are not suppressed.

history—Display entries that have already been withdrawn, but have been logged.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

suppressed—Display entries that have been suppressed and are no longer being installed into the forwarding table or exported by routing protocols.

Required Privilege Level
view

Related Documentation
- **clear bgp damping on page 1755**
- **show policy damping on page 1886**

List of Sample Output
show route damping decayed detail on page 1925
show route damping history on page 1925
show route damping history detail on page 1926

Output Fields
Table 240 on page 1922 lists the output fields for the **show route damping** command. Output fields are listed in the approximate order in which they appear.

Table 240: show route damping Output Fields

Field Name	Field Description	Level of Output
routing-table-name	Name of the routing table—for example,**inet.0**.	All levels
destinations	Number of destinations for which there are routes in the routing table.	All levels

Table 240: show route damping Output Fields *(continued)*

Field Name	Field Description	Level of Output
***number* routes**	Number of routes in the routing table and total number of routes in the following states: • **active** • **holddown** (routes that are in a pending state before being declared inactive) • **hidden** (the routes are not used because of a routing policy)	All levels
***destination-prefix* (entry, announced)**	Destination prefix. The **entry** value is the number of routes for this destination, and the **announced** value is the number of routes being announced for this destination.	**detail extensive**
[*protocol*, *preference*]	Protocol from which the route was learned and the preference value for the route. • **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • **-**—A hyphen indicates the last active route. • *****—An asterisk indicates that the route is both the active and the last active route. An asterisk before a **to** line indicates the best subpath to the route. In every routing metric except for the BGP **LocalPref** attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the **LocalPref** value in the **Preference2** field. For example, if the LocalPref value for Route 1 is 100, the Preference2 value is -101. If the **LocalPref** value for Route 2 is 155, the **Preference2** value is -156. Route 2 is preferred because it has a higher **LocalPref** value and a lower **Preference2** value.	All levels
Next-hop reference count	Number of references made to the next hop.	**detail extensive**
Source	IP address of the route source.	**detail extensive**
Next hop	Network layer address of the directly reachable neighboring system.	**detail extensive**
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word **Selected**.	**detail extensive**
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.	**detail extensive**
Indirect next hop	Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.	**detail extensive**
State	Flags for this route. For a description of possible values for this field, see the output field table for the **show route detail** command.	**detail extensive**
Local AS	AS number of the local routing device.	**detail extensive**
Peer AS	AS number of the peer routing device.	**detail extensive**

Table 240: show route damping Output Fields *(continued)*

Field Name	Field Description	Level of Output
Age	How long the route has been known.	**detail extensive**
Metric	Metric for the route.	**detail extensive**
Task	Name of the protocol that has added the route.	**detail extensive**
Announcement bits	List of protocols that announce this route. ***n*-Resolve inet** indicates that the route is used for route resolution for next hops found in the routing table. ***n*** is an index used by Dell Support only (see "Requesting Technical Support" on page lxxi).	**detail extensive**
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device or if AS path prepending is configured. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set.	All levels
to	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.	**brief** none
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word **Selected**.	**brief** none
Communities	Community path attribute for the route. See the output field table for the **show route detail** command.	**detail extensive**
Localpref	Local preference value included in the route.	All levels
Router ID	BGP router ID as advertised by the neighbor in the open message.	**detail extensive**
Merit (last update/now)	Last updated and current figure-of-merit value.	**detail extensive**
damping-parameters	Name that identifies the damping parameters used, which is defined in the damping statement at the **[edit policy-options]** hierarchy level.	**detail extensive**
Last update	Time of most recent change in path attributes.	**detail extensive**

Table 240: show route damping Output Fields *(continued)*

Field Name	Field Description	Level of Output
First update	Time of first change in path attributes, which started the route damping process.	**detail extensive**
Flaps	Number of times the route has gone up or down or its path attributes have changed.	**detail extensive**
Suppressed	(**suppressed** keyword only) This route is currently suppressed. A suppressed route does not appear in the forwarding table and routing protocols do not export it.	All levels
Reusable in	(**suppressed** keyword only) Time when a suppressed route will again be available.	All levels
Preference will be	(**suppressed** keyword only) Preference value that will be applied to the route when it is again active.	All levels

show route damping decayed detail

```
user@host> show route damping decayed detail
inet.0: 173319 destinations, 1533668 routes (172625 active, 4 holddown, 108083
hidden)
10.0.111.0/24 (7 entries, 1 announced)
        *BGP    Preference: 170/-101
                Next-hop reference count: 151973
                Source: 172.23.2.129
                Next hop: via so-1/2/0.0
                Next hop: via so-5/1/0.0, selected
                Next hop: via so-6/0/0.0
                Protocol next hop: 172.23.2.129
                Indirect next hop: 89a1a00 264185
                State: <Active Ext>
                Local AS: 65000 Peer AS:   65490
                Age: 3:28       Metric2: 0
                Task: BGP_65490.172.23.2.129+179
                Announcement bits (6): 0-KRT 1-RT 4-KRT 5-BGP.0.0.0.0+179

                6-Resolve tree 2 7-Resolve tree 3
                AS path: 65490 65520 65525 65525 65525 65525 I ()
                Communities: 65501:390 65501:2000 65501:3000 65504:701
                Localpref: 100
                Router ID: 172.23.2.129
                Merit (last update/now): 1934/1790
                damping-parameters: damping-high
                Last update:       00:03:28 First update:       00:06:40
                Flaps: 2
```

show route damping history

```
user@host> show route damping history
inet.0: 173320 destinations, 1533529 routes (172624 active, 6 holddown, 108122
hidden)
+ = Active Route, - = Last Active, * = Both

10.108.0.0/15       [BGP ] 2d 22:47:58, localpref 100
                      AS path: 65220 65501 65502 I
                    > to 192.168.60.85 via so-3/1/0.0
```

show route damping history detail

```
user@host> show route damping history detail
inet.0: 173319 destinations, 1533435 routes (172627 active, 2 holddown, 108105
hidden)
10.108.0.0/15 (3 entries, 1 announced)
         BGP                 /-101
                Next-hop reference count: 69058
                Source: 192.168.60.85
                Next hop: 192.168.60.85 via so-3/1/0.0, selected
                State: <Hidden Ext>
                Inactive reason: Unusable path
                Local AS: 65000 Peer AS:  65220
                Age: 2d 22:48:10
                Task: BGP_65220.192.168.60.85+179
                AS path: 65220 65501 65502 I ()
                Communities: 65501:390 65501:2000 65501:3000 65504:3561
                Localpref: 100
                Router ID: 192.168.80.25
                Merit (last update/now): 1000/932
                damping-parameters: set-normal
                Last update:       00:01:05 First update:       00:01:05
                Flaps: 1
```

show route detail

Syntax
show route detail
<*destination-prefix*>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show route detail
<*destination-prefix*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display detailed information about the active entries in the routing tables.

Options
none—Display all active entries in the routing table on all systems.

destination-prefix—(Optional) Display active entries for the specified address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route detail on page 1935

Output Fields
Table 241 on page 1927 describes the output fields for the **show route detail** command. Output fields are listed in the approximate order in which they appear.

Table 241: show route detail Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, **inet.0**).
number **destinations**	Number of destinations for which there are routes in the routing table.
number **routes**	Number of routes in the routing table and total number of routes in the following states: • **active** (routes that are active) • **holddown** (routes that are in the pending state before being declared inactive) • **hidden** (routes that are not used because of a routing policy)

Table 241: show route detail Output Fields *(continued)*

Field Name	Field Description
route-destination **(entry, announced)**	Route destination (for example:**10.0.0.1/24**). The **entry** value is the number of routes for this destination, and the **announced** value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as: • ***MPLS-label*** (for example, **80001**). • ***interface-name*** (for example, **ge-1/0/2**). • ***neighbor-address*:*control-word-status*:*encapsulation type*:*vc-id*:*source*** (Layer 2 circuit only; for example, **10.1.1.195:NoCtrlWord:1:1:Local/96**). • ***neighbor-address***—Address of the neighbor. • ***control-word-status***—Whether the use of the control word has been negotiated for this virtual circuit: **NoCtrlWord** or **CtrlWord**. • ***encapsulation type***—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport • ***vc-id***—Virtual circuit identifier. • ***source***—Source of the advertisement: **Local** or **Remote**.
label stacking	**(**Next-to-the-last-hop routing device for MPLS only) Depth of the Multiprotocol Label Switching (MPLS) label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels. • **S=0 route** indicates that a packet with an incoming label stack depth of 2 or more exits this routing device with one fewer label (the label-popping operation is performed). • If there is no **S=** information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).
[*protocol*, *preference*]	Protocol from which the route was learned and the preference value for the route. • **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • **-** —A hyphen indicates the last active route. • *****—An asterisk indicates that the route is both the active and the last active route. An asterisk before a **to** line indicates the best subpath to the route. In every routing metric except for the BGP **LocalPref** attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the **LocalPref** value in the **Preference2** field. For example, if the **LocalPref** value for Route 1 is 100, the **Preference2** value is -101. If the **LocalPref** value for Route 2 is 155, the **Preference2** value is -156. Route 2 is preferred because it has a higher **LocalPref** value and a lower **Preference2** value.
Level	(IS-IS only). In IS-IS, a single AS can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area; when the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.
Route Distinguisher	IP subnet augmented with a 64-bit prefix.
Next-hop type	Type of next hop. For a description of possible values for this field, see Table 242 on page 1931.

Table 241: show route detail Output Fields *(continued)*

Field Name	Field Description
Next-hop reference count	Number of references made to the next hop.
Source	IP address of the route source.
Next hop	Network layer address of the directly reachable neighboring system.
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of interface that is actually used is followed by the word **Selected**. This field can also contain the following information: • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when Multiprotocol Label Switching (MPLS) label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable Border Gateway Protocol (BGP) multipath load balancing.
Label-switched-path *lsp-path-name*	Name of the label-switched path (LSP) used to reach the next hop.
Label operation	MPLS label and operation occurring at this routing device. The operation can be **pop** (where a label is removed from the top of the stack), **push** (where another label is added to the label stack), or **swap** (where a label is replaced by another label).
Interface	(Local only) Local interface name.
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.
Indirect next hop	Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.
State	State of the route (a route can be in more than one state). See Table 243 on page 1932.
Local AS	AS number of the local routing device.
Age	How long the route has been known.
Metric*n*	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
MED-plus-IGP	Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.
Task	Name of the protocol that has added the route.

Table 241: show route detail Output Fields *(continued)*

Field Name	Field Description
Announcement bits	List of protocols that announce this route. ***n*-Resolve inet** indicates that the route is used for route resolution for next hops found in the routing table. ***n*** is an index used by Dell Support only (see "Requesting Technical Support" on page lxxi).
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used in the AS-path merge process, as defined in RFC 4893. • **[ ]**—If more than one AS number is configured on the routing device, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set.
VC Label	MPLS label assigned to the Layer 2 circuit virtual connection.
MTU	Maximum transmission unit (MTU) of the Layer 2 circuit.
VLAN ID	VLAN identifier of the Layer 2 circuit.
Prefixes bound to route	Forwarding Equivalent Class (FEC) bound to this route. Applicable only to routes installed by LDP.
Communities	Community path attribute for the route. See Table 244 on page 1934 for all possible values for this field.
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).
control flags	Control flags: **none** or **Site Down**.
mtu	Maximum transmission unit (MTU) information.
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.
status vector	Layer 2 VPN and VPLS network layer reachability information (NLRI).
Localpref	Local preference value included in the route.
Router ID	BGP router ID as advertised by the neighbor in the open message.

Table 241: show route detail Output Fields *(continued)*

Field Name	Field Description
Primary Routing Table	In a routing table group, the name of the primary routing table in which the route resides.
Secondary Tables	In a routing table group, the name of one or more secondary tables in which the route resides.

Table 242 on page 1931 describes all possible values for the **Next-hop Types** output field.

Table 242: Next-Hop Types Output Field Values

Next-Hop Type	Description
broadcast (bcast)	Broadcast next hop.
deny	Deny next hop.
hold	Next hop is waiting to be resolved into a unicast or multicast type.
indexed (idxd)	Indexed next hop.
indirect (indr)	Indirect next hop.
local (locl)	Local address on an interface.
routed multicast (mcrt)	Regular multicast next hop.
multicast (mcst)	Wire multicast next hop (limited to the LAN).
multicast discard (mdsc)	Multicast discard.
multicast group (mgrp)	Multicast group member.
receive (recv)	Receive.
reject (rjct)	Discard. An ICMP unreachable message was sent.
resolve (rslv)	Resolving next hop.
unicast (ucst)	Unicast.
unilist (ulst)	List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.

Table 243 on page 1932 describes all possible values for the **State** output field. A route can be in more than one state (for example, **<Active NoReadvrt Int Ext>**).

Table 243: State Output Field Values

Value	Description
Accounting	Route needs accounting.
Active	Route is active.
Always Compare MED	Path with a lower multiple exit discriminator (MED) is available.
AS path	Shorter AS path is available.
Clone	Route is a clone.
Cisco Non-deterministic MED selection	Cisco nondeterministic MED is enabled and a path with a lower MED is available.
Cluster list length	Length of cluster list sent by the route reflector.
Delete	Route has been deleted.
Ex	Exterior route.
Ext	BGP route received from an external BGP neighbor.
FlashAll	Forces all protocols to be notified of a change to any route, active or inactive, for a prefix. When not set, protocols are informed of a prefix only when the active route changes.
Hidden	Route not used because of routing policy.
IfCheck	Route needs forwarding RPF check.
IGP metric	Path through next hop with lower IGP metric is available.
Local Preference	Path with a higher local preference value is available.
Inactive reason	Flags for this route, which was not selected as best for a particular destination.
Initial	Route being added.
Int	Interior route.
Int Ext	BGP route received from an internal BGP peer or a BGP confederation peer.

Table 243: State Output Field Values *(continued)*

Value	Description
Interior > Exterior > Exterior via Interior	Direct, static, IGP, or EBGP path is available.
Martian	Route is a martian (ignored because it is obviously invalid).
MartianOK	Route exempt from martian filtering.
Next hop address	Path with lower metric next hop is available.
No difference	Path from neighbor with lower IP address is available.
NoReadvrt	Route not to be advertised.
NotBest	Route not chosen because it does not have the lowest MED.
Not Best in its group	Incoming BGP AS is not the best of a group (only one AS can be the best).
NotInstall	Route not to be installed in the forwarding table.
Number of gateways	Path with greater number of next hops is available.
Origin	Path with lower origin code is available.
Pending	Route pending because of a hold-down configured on another route.
Release	Route scheduled for release.
RIB preference	Route from a higher-numbered routing table is available.
Route Distinguisher	64-bit prefix added to IP subnets to make them unique.
Route Metric or MED comparison	Route with a lower metric or MED is available.
Route Preference	Route with lower preference value is available
Router ID	Path through neighbor with lower ID is available.
Secondary	Route not a primary route.
Unusable path	Path is not usable because of one of the following conditions: • The route is damped. • The route is rejected by an import policy. • The route is unresolved.
Update source	Last tiebreaker is the lowest IP address value.

Table 244 on page 1934 describes the possible values for the **Communities** output field.

Table 244: Communities Output Field Values

Value	Description
area-number	4 bytes, encoding a 32-bit area number. For AS-external routes, the value is **0**. A nonzero value identifies the route as internal to the OSPF domain, and as within the identified area. Area numbers are relative to a particular OSPF domain.
bandwidth: *local AS number:link-bandwidth-number*	Link-bandwidth community value used for unequal-cost load balancing. When BGP has several candidate paths available for multipath purposes, it does not perform unequal-cost load balancing according to the link-bandwidth community unless all candidate paths have this attribute.
domain-id	Unique configurable number that identifies the OSPF domain.
domain-id-vendor	Unique configurable number that identifies the OSPF domain.
link-bandwidth-number	Link-bandwidth number: from **0** through **4,294,967,295** (bytes per second).
local AS number	Local AS number: from **1** through **65,535**.
options	1 byte. Currently this is only used if the route type is **5** or **7**. Setting the least significant bit in the field indicates that the route carries a type 2 metric.
origin	(Used with VPNs) Identifies where the route came from.
ospf-route-type	1 byte, encoded as **1** or **2** for intra-area routes (depending on whether the route came from a type 1 or a type 2 LSA); **3** for summary routes; **5** for external routes (area number must be **0**); **7** for NSSA routes; or **129** for sham link endpoint addresses.
rte-type	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute **0x0306**. The format is *area-number:ospf-route-type:options*.
route-type-vendor	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute **0x8000**. The format is *area-number:ospf-route-type:options*.
target	Defines which VPN the route participates in; **target** has the format *32-bit IP address:16-bit number*. For example, **10.19.0.0:100**.
unknown IANA	Incoming IANA codes with a value between **0x1** and **0x7fff**. This code of the BGP extended community attribute is accepted, but it is not recognized.
unknown OSPF vendor community	Incoming IANA codes with a value above **0x8000**. This code of the BGP extended community attribute is accepted, but it is not recognized.

show route detail `user@host> show route detail`

```
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 29
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 1:31:43
                Task: RT
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

10.31.1.0/30 (2 entries, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 2
                Next hop: via so-0/3/0.0, selected
                State: <Active Int>
                Local AS:    69
                Age: 1:30:17
                Task: IF
                Announcement bits (1): 3-Resolve tree 2
                AS path: I
         OSPF   Preference: 10
                Next-hop reference count: 1
                Next hop: via so-0/3/0.0, selected
                State: <Int>
                Inactive reason: Route Preference
                Local AS:    69
                Age: 1:30:17    Metric: 1
                Area: 0.0.0.0
                Task: OSPF
                AS path: I

10.31.1.1/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 7
                Interface: so-0/3/0.0
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:30:20
                Task: IF
                Announcement bits (1): 3-Resolve tree 2
                AS path: I

...
```

```
10.31.2.0/30 (1 entry, 1 announced)
        *OSPF   Preference: 10
                Next-hop reference count: 9
                Next hop: via so-0/3/0.0
                Next hop: 10.31.1.6 via ge-3/1/0.0, selected
                State: <Active Int>
                Local AS:    69
                Age: 1:29:56    Metric: 2
                Area: 0.0.0.0
                Task: OSPF
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

...

224.0.0.2/32 (1 entry, 1 announced)
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:45
                Task: PIM Recv
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

...

224.0.0.22/32 (1 entry, 1 announced)
        *IGMP   Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:43
                Task: IGMP
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.255.70.103/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 6
                Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 100096
                State: <Active Int>
                Local AS:    69
                Age: 1:25:49    Metric: 2
                Task: RSVP
                Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
                AS path: I

10.255.71.238/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 6
                Next hop: via so-0/3/0.0 weight 0x1, selected
                Label-switched-path green-r1-r2
                State: <Active Int>
                Local AS:    69
```

```
                Age: 1:25:49    Metric: 1
                Task: RSVP
                Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
                AS path: I

private__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Local AS:    69
                Age: 1:31:44
                Task: IF
                AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
0   (1 entry, 1 announced)
        *MPLS   Preference: 0
                Next hop type: Receive
                Next-hop reference count: 6
                State: <Active Int>
                Local AS:    69
                Age: 1:31:45    Metric: 1
                Task: MPLS
                Announcement bits (1): 0-KRT
                AS path: I

...

800010 (1 entry, 1 announced)
        *VPLS   Preference: 7
                Next-hop reference count: 2
                Next hop: via vt-3/2/0.32769, selected
                Label operation: Pop
                State: <Active Int>
                Age: 1:29:30
                Task: Common L2 VC
                Announcement bits (1): 0-KRT
                AS path: I

vt-3/2/0.32769 (1 entry, 1 announced)
        *VPLS   Preference: 7
                Next-hop reference count: 2
                Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 800012, Push 100096(top)
                Protocol next hop: 10.255.70.103
                Push 800012
                Indirect next hop: 87272e4 1048574
                State: <Active Int>
                Age: 1:29:30    Metric2: 2
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 1-Common L2 VC
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0
```

```
inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

abcd::10:255:71:52/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Local AS:    69
                Age: 1:31:44
                Task: IF
                AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:44
                Task: IF
                AS path: I

ff02::2/128 (1 entry, 1 announced)
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:45
                Task: PIM Recv6
                Announcement bits (1): 0-KRT
                AS path: I

ff02::d/128 (1 entry, 1 announced)
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:45
                Task: PIM Recv6
                Announcement bits (1): 0-KRT
                AS path: I

ff02::16/128 (1 entry, 1 announced)
        *MLD    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:43
                Task: MLD
                Announcement bits (1): 0-KRT
                AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
```

```
                Next hop: via lo0.16385, selected
                State: <Active NoReadvrt Int>
                Age: 1:31:44
                Task: IF
                AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

10.255.70.103:1:3:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.70.103:1
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 1:25:49    Metric2: 1
                Task: BGP_69.10.255.70.103+179
                Announcement bits (1): 0-green-l2vpn
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0
                Label-base: 800008, range: 8
                Localpref: 100
                Router ID: 10.255.70.103
                Primary Routing Table bgp.l2vpn.0

10.255.71.52:1:1:1/96 (1 entry, 1 announced)
        *L2VPN  Preference: 170/-1
                Next-hop reference count: 5
                Protocol next hop: 10.255.71.52
                Indirect next hop: 0 -
                State: <Active Int Ext>
                Age: 1:31:40    Metric2: 1
                Task: green-l2vpn
                Announcement bits (1): 1-BGP.0.0.0.0+179
                AS path: I
                Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
                mtu: 0
                Label-base: 800016, range: 8, status-vector: 0x9F

10.255.71.52:1:5:1/96 (1 entry, 1 announced)
        *L2VPN  Preference: 170/-101
                Next-hop reference count: 5
                Protocol next hop: 10.255.71.52
                Indirect next hop: 0 -
                State: <Active Int Ext>
                Age: 1:31:40    Metric2: 1
                Task: green-l2vpn
                Announcement bits (1): 1-BGP.0.0.0.0+179
                AS path: I
                Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
                Label-base: 800008, range: 8, status-vector: 0x9F

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
        *L2CKT  Preference: 7
                Next hop: via so-1/1/2.0 weight 1, selected
```

```
Label-switched-path my-lsp
Label operation: Push 100000[0]
Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
State: <Active Int>
Local AS:    99
Age: 10:21
Task: l2 circuit
Announcement bits (1): 0-LDP
AS path: I
VC Label 100000, MTU 1500, VLAN ID 512
```

show route exact

Syntax
show route exact *destination-prefix*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show route exact *destination-prefix*
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display only the routes that exactly match the specified address or range of addresses.

Options
brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.

destination-prefix—Address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route exact on page 1941
show route exact detail on page 1941
show route exact extensive on page 1943
show route exact terse on page 1943

Output Fields
For information about output fields, see the output field tables for the **show route** command, the **show route detail** command, the **show route extensive** command, or the **show route terse** command.

show route exact

```
user@host> show route exact 207.17.136.0/24

inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
207.17.136.0/24    *[Static/5] 2d 03:30:22
                     > to 192.168.71.254 via fxp0.0
```

show route exact detail

```
user@host> show route exact 207.17.136.0/24 detail

inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
207.17.136.0/24 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 29
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2d 3:30:26
                Task: RT
```

```
                    Announcement bits (2): 0-KRT 3-Resolve tree 2
                    AS path: I
```

show route exact extensive

```
user@host> show route exact 207.17.136.0/24 extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
207.17.136.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.0/24 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 29
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 1:25:18
                Task: RT
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I
```

show route exact terse

```
user@host> show route exact 207.17.136.0/24 terse

inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 207.17.136.0/24    S   5                      >192.168.71.254
```

show route export

Syntax
show route export
<brief | detail>
<instance <*instance-name*> | *routing-table-name*>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show route export
<brief | detail>
<instance <*instance-name*> | *routing-table-name*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display policy-based route export information. Policy-based export simplifies the process of exchanging route information between routing instances.

Options
none—(Same as **brief**.) Display standard information about policy-based export for all instances and routing tables on all systems.

brief | detail—(Optional) Display the specified level of output.

instance <*instance-name*>—(Optional) Display a particular routing instance for which policy-based export is currently enabled.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

routing-table-name—(Optional) Display information about a particular routing table (for example, **inet.0**) for which policy-based export is currently enabled. (For information about the different types of routing tables, see the *Junos OS Routing Protocols Configuration Guide*.)

Required Privilege Level
view

List of Sample Output
show route export on page 1945
show route export detail on page 1945
show route export instance detail on page 1945

Output Fields
Table 245 on page 1944 lists the output fields for the **show route export** command. Output fields are listed in the approximate order in which they appear.

Table 245: show route export Output Fields

Field Name	Field Description	Level of Output
Table or ***table-name***	Name of the routing tables that either import or export routes.	All levels
Routes	Number of routes exported from this table into other tables. If a particular route is exported to different tables, the counter will only increment by one.	**brief** none
Export	Whether the table is currently exporting routes to other tables: **Y** or **N** (Yes or No).	**brief** none

Table 245: show route export Output Fields *(continued)*

Field Name	Field Description	Level of Output
Import	Tables currently importing routes from the originator table. (Not displayed for tables that are not exporting any routes.)	**detail**
Flags	(**instance** keyword only) Flags for this feature on this instance: • **config auto-policy**—The policy was deduced from the configured IGP export policies. • **cleanup**—Configuration information for this instance is no longer valid. • **config**—The instance was explicitly configured.	**detail**
Options	(**instance** keyword only) Configured option displays the type of routing tables the feature handles: • **unicast**—Indicates ***instance*.inet.0**. • **multicast**—Indicates ***instance*.inet.2**. • **unicast multicast**—Indicates ***instance*.inet.0** and ***instance*.inet.2**.	**detail**
Import policy	(**instance** keyword only) Policy that **route export** uses to construct the import-export matrix. Not displayed if the instance type is **vrf**.	**detail**
Instance	(**instance** keyword only) Name of the routing instance.	**detail**
Type	(**instance** keyword only) Type of routing instance: **forwarding**, **non-forwarding**, or **vrf**.	**detail**

show route export

```
user@host> show route export
Table                            Export           Routes
inet.0                           N                     0
black.inet.0                     Y                     3
red.inet.0                       Y                     4
```

show route export detail

```
user@host> show route export detail
inet.0                                           Routes:        0
black.inet.0                                     Routes:        3
  Import: [ inet.0 ]
red.inet.0                                       Routes:        4
  Import: [ inet.0 ]
```

show route export instance detail

```
user@host> show route export instance detail
Instance: master                         Type: forwarding
  Flags: <config auto-policy> Options: <unicast multicast>
  Import policy: [ (ospf-master-from-red || isis-master-from-black) ]
Instance: black                          Type: non-forwarding
Instance: red                            Type: non-forwarding
```

show route extensive

Syntax
show route extensive
<*destination-prefix*>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show route extensive
<*destination-prefix*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display extensive information about the active entries in the routing tables.

Options
none—Display all active entries in the routing table.

destination-prefix—(Optional) Display active entries for the specified address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route extensive on page 1951
show route extensive (Access Route) on page 1957
show route extensive (Route Reflector) on page 1957

Output Fields
Table 246 on page 1946 describes the output fields for the **show route extensive** command. Output fields are listed in the approximate order in which they appear.

Table 246: show route extensive Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, **inet.0**).
***number* destinations**	Number of destinations for which there are routes in the routing table.
***number* routes**	Number of routes in the routing table and total number of routes in the following states: • **active** (routes that are active). • **holddown** (routes that are in the pending state before being declared inactive). • **hidden** (routes that are not used because of a routing policy).

Table 246: show route extensive Output Fields *(continued)*

Field Name	Field Description
route-destination **(entry, announced)**	Route destination (for example:**10.0.0.1/24**). The **entry** value is the number of route for this destination, and the **announced** value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as: • ***MPLS-label*** (for example, **80001**). • ***interface-name*** (for example, **ge-1/0/2**). • ***neighbor-address*****:*****control-word-status*****:*****encapsulation type*****:*****vc-id*****:*****source*** (Layer 2 circuit only; for example, **10.1.1.195:NoCtrlWord:1:1:Local/96**). • ***neighbor-address***—Address of the neighbor. • ***control-word-status***—Whether the use of the control word has been negotiated for this virtual circuit: **NoCtrlWord** or **CtrlWord**. • ***encapsulation type***—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport. • ***vc-id***—Virtual circuit identifier. • ***source***—Source of the advertisement: **Local** or **Remote**.
TSI	Protocol header information.
label stacking	(Next-to-the-last-hop routing device for MPLS only) Depth of the Multiprotocol Label Switching (MPLS) label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels. • **S=0 route** indicates that a packet with an incoming label stack depth of two or more exits this router with one fewer label (the label-popping operation is performed). • If there is no **S=** information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).
[*protocol*, *preference*]	Protocol from which the route was learned and the preference value for the route. • **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • **-** —A hyphen indicates the last active route. • *****—An asterisk indicates that the route is both the active and the last active route. An asterisk before a **to** line indicates the best subpath to the route. In every routing metric except for the BGP **LocalPref** attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the **LocalPref** value in the **Preference2** field. For example, if the **LocalPref** value for Route 1 is 100, the **Preference2** value is -101. If the **LocalPref** value for Route 2 is 155, the **Preference2** value is -156. Route 2 is preferred because it has a higher **LocalPref** value and a lower **Preference2** value.
Level	(IS-IS only). In IS-IS, a single autonomous system (AS) can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area; when the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.

Table 246: show route extensive Output Fields *(continued)*

Field Name	Field Description
Route Distinguisher	IP subnet augmented with a 64-bit prefix.
Next-hop type	Type of next hop. For a description of possible values for this field, see the Output Field table in the **show route detail** command.
Next-hop reference count	Number of references made to the next hop.
Source	IP address of the route source.
Next hop	Network layer address of the directly reachable neighboring system.
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word **Selected**. This field can also contain the following information: • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when Multiprotocol Label Switching (MPLS) label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable Border Gateway Protocol (BGP) multipath load balancing.
Label-switched-path *lsp-path-name*	Name of the label-switched path (LSP) used to reach the next hop.
Label operation	MPLS label and operation occurring at this routing device. The operation can be **pop** (where a label is removed from the top of the stack), **push** (where another label is added to the label stack), or **swap** (where a label is replaced by another label).
Offset	Whether the metric has been increased or decreased by an offset value.
Interface	(Local only) Local interface name.
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to recursively derive a forwarding next hop.
label-operation	MPLS label and operation occurring at this routing device. The operation can be **pop** (where a label is removed from the top of the stack), **push** (where another label is added to the label stack), or **swap** (where a label is replaced by another label).
Indirect next hops	When present, a list of nodes that are used to resolve the path to the next-hop destination, in the order that they are resolved.
State	State of the route (a route can be in more than one state). See the Output Field table in the **show route detail** command.

Table 246: show route extensive Output Fields *(continued)*

Field Name	Field Description
Inactive reason	If the route is inactive, the reason for its current state is indicated. Typical reasons include: • **Active preferred**—Currently active route was selected over this route. • **Always compare MED**—Path with a lower multiple exit discriminator (MED) is available. • **AS path**—Shorter AS path is available. • **Cisco Non-deterministic MED selection**—Cisco nondeterministic MED is enabled and a path with a lower MED is available. • **Cluster list length**—Path with a shorter cluster list length is available. • **Forwarding use only**—Path is only available for forwarding purposes. • **IGP metric**—Path through the next hop with a lower IGP metric is available. • **IGP metric type**—Path with a lower OSPF link-state advertisement type is available. • **Interior > Exterior > Exterior via Interior**—Direct, static, IGP, or EBGP path is available. • **Local preference**—Path with a higher local preference value is available. • **Next hop address**—Path with a lower metric next hop is available. • **No difference**—Path from a neighbor with a lower IP address is available. • **Not Best in its group**—Occurs when multiple peers of the same external AS advertise the same prefix and are grouped together in the selection process. When this reason is displayed, an additional reason is provided (typically one of the other reasons listed). • **Number of gateways**—Path with a higher number of next hops is available. • **Origin**—Path with a lower origin code is available. • **OSPF version**—Path does not support the indicated OSPF version. • **RIB preference**—Route from a higher-numbered routing table is available. • **Route destinguisher**—64-bit prefix added to IP subnets to make them unique. • **Route metric or MED comparison**—Route with a lower metric or MED is available. • **Route preference**—Route with a lower preference value is available. • **Router ID**—Path through a neighbor with a lower ID is available. • **Unusable path**—Path is not usable because of one of the following conditions: the route is damped, the route is rejected by an import policy, or the route is unresolved. • **Update source**—Last tiebreaker is the lowest IP address value.
Local AS	Autonomous system (AS) number of the local routing device.
Age	How long the route has been known.
Metric	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
MED-plus-IGP	Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.
Task	Name of the protocol that has added the route.
Announcement bits	List of protocols that announce this route. ***n*-Resolve inet** indicates that the route is used for route resolution for next hops found in the routing table. ***n*** is an index used by Dell Support only (see “Requesting Technical Support” on page lxxi).

Table 246: show route extensive Output Fields *(continued)*

Field Name	Field Description
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set.
AS path: I <Originator>	(For route reflected output only) Originator ID attribute set by the route reflector.
VC Label	MPLS label assigned to the Layer 2 circuit virtual connection.
MTU	Maximum transmission unit (MTU) of the Layer 2 circuit.
VLAN ID	VLAN identifier of the Layer 2 circuit.
Cluster list	(For route reflected output only) Cluster ID sent by the route reflector.
Originator ID	(For route reflected output only) Address of router that originally sent the route to the route reflector.
Prefixes bound to route	Forwarding Equivalent Class (FEC) bound to this route. Applicable only to routes installed by LDP.
Communities	Community path attribute for the route. See the Output Field table in the **show route detail** command for all possible values for this field.
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).
control flags	Control flags: **none** or Site Down.
mtu	Maximum transmission unit (MTU) information.
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.
status vector	Layer 2 VPN and VPLS network layer reachability information (NLRI).
Localpref	Local preference value included in the route.
Router ID	BGP router ID as advertised by the neighbor in the open message.

Table 246: show route extensive Output Fields *(continued)*

Field Name	Field Description
Primary Routing Table	In a routing table group, the name of the primary routing table in which the route resides.
Secondary Tables	In a routing table group, the name of one or more secondary tables in which the route resides.
Originating RIB	Name of the routing table whose active route was used to determine the forwarding next-hop entry in the resolution database. For example, in the case of **inet.0** resolving through **inet.0** and **inet.3**, this field indicates which routing table, **inet.0** or **inet.3**, provided the best path for a particular prefix.
Node path count	Number of nodes in the path.
Forwarding nexthops	Number of forwarding next hops. The forwarding next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.

show route extensive

```
user@host> show route extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 29
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 1:34:06
                Task: RT
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

10.31.1.0/30 (2 entries, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 2
                Next hop: via so-0/3/0.0, selected
                State: <Active Int>
                Local AS:    69
                Age: 1:32:40
                Task: IF
                Announcement bits (1): 3-Resolve tree 2
                AS path: I
```

```
         OSPF   Preference: 10
                Next-hop reference count: 1
                Next hop: via so-0/3/0.0, selected
                State: <Int>
                Inactive reason: Route Preference
                Local AS:    69
                Age: 1:32:40    Metric: 1
                Area: 0.0.0.0
                Task: OSPF
                AS path: I

10.31.1.1/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 7
                Interface: so-0/3/0.0
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:32:43
                Task: IF
                Announcement bits (1): 3-Resolve tree 2
                AS path: I

...

10.31.2.0/30 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.31.2.0/30 -> {10.31.1.6}
        *OSPF   Preference: 10
                Next-hop reference count: 9
                Next hop: via so-0/3/0.0
                Next hop: 10.31.1.6 via ge-3/1/0.0, selected
                State: <Active Int>
                Local AS:    69
                Age: 1:32:19    Metric: 2
                Area: 0.0.0.0
                Task: OSPF
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

...

224.0.0.2/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 224.0.0.2/32 -> {}
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:34:08
                Task: PIM Recv
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

...

224.0.0.22/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 224.0.0.22/32 -> {}
        *IGMP   Preference: 0
                Next-hop reference count: 18
```

```
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:34:06
                Task: IGMP
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.255.70.103/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 6
                Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 100096
                State: <Active Int>
                Local AS:    69
                Age: 1:28:12    Metric: 2
                Task: RSVP
                Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
                AS path: I

10.255.71.238/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 6
                Next hop: via so-0/3/0.0 weight 0x1, selected
                Label-switched-path green-r1-r2
                State: <Active Int>
                Local AS:    69
                Age: 1:28:12    Metric: 1
                Task: RSVP
                Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

...

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Local AS:    69
                Age: 1:34:07
                Task: IF
                AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

0   (1 entry, 1 announced)
TSI:
KRT in-kernel 0      /36 -> {}
        *MPLS   Preference: 0
                Next hop type: Receive
                Next-hop reference count: 6
```

```
                State: <Active Int>
                Local AS:    69
                Age: 1:34:08    Metric: 1
                Task: MPLS
                Announcement bits (1): 0-KRT
                AS path: I

...

800010 (1 entry, 1 announced)

TSI:
KRT in-kernel 800010 /36 -> {vt-3/2/0.32769}
        *VPLS   Preference: 7
                Next-hop reference count: 2
                Next hop: via vt-3/2/0.32769, selected
                Label operation: Pop
                State: <Active Int>
                Age: 1:31:53
                Task: Common L2 VC
                Announcement bits (1): 0-KRT
                AS path: I

vt-3/2/0.32769 (1 entry, 1 announced)
TSI:
KRT in-kernel vt-3/2/0.32769.0      /16 -> {indirect(1048574)}
        *VPLS   Preference: 7
                Next-hop reference count: 2
                Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 800012, Push 100096(top)
                Protocol next hop: 10.255.70.103
                Push 800012
                Indirect next hop: 87272e4 1048574
                State: <Active Int>
                Age: 1:31:53    Metric2: 2
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 1-Common L2 VC
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0
                Indirect next hops: 1
                        Protocol next hop: 10.255.70.103 Metric: 2
                        Push 800012
                        Indirect next hop: 87272e4 1048574
                        Indirect path forwarding next hops: 1
                                Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1
                        10.255.70.103/32 Originating RIB: inet.3
                          Metric: 2                     Node path count: 1
                          Forwarding nexthops: 1
                                Nexthop: 10.31.1.6 via ge-3/1/0.0

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

abcd::10:255:71:52/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Local AS:    69
```

```
                Age: 1:34:07
                Task: IF
                AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:34:07
                Task: IF
                AS path: I

ff02::2/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::2/128 -> {}
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:34:08
                Task: PIM Recv6
                Announcement bits (1): 0-KRT
                AS path: I

ff02::d/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::d/128 -> {}
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:34:08
                Task: PIM Recv6
                Announcement bits (1): 0-KRT
                AS path: I

ff02::16/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::16/128 -> {}
        *MLD    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:34:06
                Task: MLD
                Announcement bits (1): 0-KRT
                AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.16385, selected
                State: <Active NoReadvrt Int>
                Age: 1:34:07
                Task: IF
```

```
                AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

10.255.70.103:1:3:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.70.103:1
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 1:28:12    Metric2: 1
                Task: BGP_69.10.255.70.103+179
                Announcement bits (1): 0-green-l2vpn
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0
                Label-base: 800008, range: 8
                Localpref: 100
                Router ID: 10.255.70.103
                Primary Routing Table bgp.l2vpn.0

10.255.71.52:1:1:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699540
        *L2VPN  Preference: 170/-1
                Next-hop reference count: 5
                Protocol next hop: 10.255.71.52
                Indirect next hop: 0 -
                State: <Active Int Ext>
                Age: 1:34:03    Metric2: 1
                Task: green-l2vpn
                Announcement bits (1): 1-BGP.0.0.0.0+179
                AS path: I
                Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
                mtu: 0
                Label-base: 800016, range: 8, status-vector: 0x9F

10.255.71.52:1:5:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699528
        *L2VPN  Preference: 170/-101
                Next-hop reference count: 5
                Protocol next hop: 10.255.71.52
                Indirect next hop: 0 -
                State: <Active Int Ext>
                Age: 1:34:03    Metric2: 1
                Task: green-l2vpn
                Announcement bits (1): 1-BGP.0.0.0.0+179
                AS path: I
                Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
                Label-base: 800008, range: 8, status-vector: 0x9F

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

TSI:
```

```
10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
        *L2CKT  Preference: 7
                Next hop: via so-1/1/2.0 weight 1, selected
                Label-switched-path my-lsp
                Label operation: Push 100000[0]
                Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
                State: <Active Int>
                Local AS:    99
                Age: 10:21
                Task: l2 circuit
                Announcement bits (1): 0-LDP
                AS path: I
                VC Label 100000, MTU 1500, VLAN ID 512
```

show route extensive (Access Route)

```
user@host> show route 13.160.0.102 extensive
inet.0: 39256 destinations, 39258 routes (39255 active, 0 holddown, 1 hidden)
13.160.0.102/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 13.160.0.102/32 -> {13.160.0.2}
OSPF area : 0.0.0.0, LSA ID : 13.160.0.102, LSA type : Extern
        *Access Preference: 13
                Next-hop reference count: 78472
                Next hop: 13.160.0.2 via fe-0/0/0.0, selected
                State: <Active Int>
       Age: 12
                Task: RPD Unix Domain Server./var/run/rpd_serv.local
                Announcement bits (2): 0-KRT 1-OSPFv2
                AS path: I
```

show route extensive (Route Reflector)

```
user@host> show route extensive
1.0.0.0/8 (1 entry, 1 announced)

TSI:
KRT in-kernel 1.0.0.0/8 -> {indirect(40)}
        *BGP    Preference: 170/-101
                Source: 192.168.4.214
                Protocol next hop: 207.17.136.192 Indirect next hop: 84ac908 40
                State: <Active Int Ext>
                Local AS: 10458 Peer AS: 10458
                Age: 3:09       Metric: 0       Metric2: 0
                Task: BGP_10458.192.168.4.214+1033
                Announcement bits (2): 0-KRT 4-Resolve inet.0
                AS path: 3944 7777 I <Originator>
                Cluster list:  1.1.1.1
                Originator ID: 10.255.245.88
                Communities: 7777:7777
                Localpref: 100
                Router ID: 4.4.4.4
                Indirect next hops: 1
                        Protocol next hop: 207.17.136.192 Metric: 0
                        Indirect next hop: 84ac908 40
                        Indirect path forwarding next hops: 0
                                Next hop type: Discard
```

show route flow validation

Syntax
show route flow validation
<brief | detail>
<*ip-prefix*>
<table *table-name*>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show route flow validation
<brief | detail>
<*ip-prefix*>
<table *table-name*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display flow route information.

Options
none—Display flow route information.

brief | detail—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.

ip-prefix—(Optional) IP address for the flow route.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

table *table-name*—(Optional) Name of the flow route table.

Required Privilege Level
view

List of Sample Output
show route flow validation on page 1959

Output Fields
Table 247 on page 1958 lists the output fields for the **show route flow validation** command. Output fields are listed in the approximate order in which they appear.

Table 247: show route flow validation Output Fields

Field Name	Field Description	Level of Output
routing-table-name	Name of the routing table (for example, **inet.0**).	All levels
prefix	Route address.	All levels
Active unicast route	Active route in the routing table.	All levels
Dependent flow destinations	Number of flows for which there are routes in the routing table.	All levels
Origin	Source of the route flow.	All levels

Table 247: show route flow validation Output Fields *(continued)*

Field Name	Field Description	Level of Output
Neighbor AS	Autonomous system identifier of the neighbor.	All levels
Flow destination	Number of entries and number of destinations that match the route flow.	All levels
Unicast best match	Destination that is the best match for the route flow.	All levels
Flags	Information about the route flow.	All levels

show route flow validation

```
user@host> show route flow validation
inet.0:
10.0.5.0/24Active unicast route
Dependent flow destinations: 1
Origin: 192.168.224.218, Neighbor AS: 65001
Flow destination (3 entries, 1 match origin)
Unicast best match: 10.0.5.0/24
Flags: SubtreeApex Consistent
```

show route inactive-path

Syntax
show route inactive-path
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show route inactive-path
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display routes for destinations that have no active route. An inactive route is a route that was not selected as the best path.

Options
none—Display all inactive routes.

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route inactive-path on page 1960
show route inactive-path detail on page 1961
show route inactive-path extensive on page 1962
show route inactive-path terse on page 1962

Output Fields
For information about output fields, see the output field tables for the **show route** command, the **show route detail** command, the **show route extensive** command, or the **show route terse** command.

show route inactive-path

```
user@host> show route inactive-path

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.12.100.12/30     [OSPF/10] 03:57:28, metric 1
                    > via so-0/3/0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.0/8          [Direct/0] 04:39:56
                    > via fxp1.0

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.12.80.0/30       [BGP/170] 04:38:17, localpref 100
```

```
                        AS path: 100 I
                      > to 10.12.80.1 via ge-6/3/2.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route inactive-path detail

```
user@host> show route inactive-path detail

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete

10.12.100.12/30 (2 entries, 1 announced)
         OSPF   Preference: 10
                Next-hop reference count: 1
                Next hop: via so-0/3/0.0, selected
                State: <Int>
                Inactive reason: Route Preference
                Local AS:     1
                Age: 3:58:24    Metric: 1
                Area: 0.0.0.0
                Task: OSPF
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

10.0.0.0/8 (2 entries, 0 announced)
         Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via fxp1.0, selected
                State: <NotBest Int>
                Inactive reason: No difference
                Age: 4:40:52
                Task: IF
                AS path: I

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete

10.12.80.0/30 (2 entries, 1 announced)
         BGP    Preference: 170/-101
                Next-hop reference count: 6
                Source: 10.12.80.1
                Next hop: 10.12.80.1 via ge-6/3/2.0, selected
                State: <Ext>
                Inactive reason: Route Preference
                Peer AS:   100
                Age: 4:39:13
                Task: BGP_100.10.12.80.1+179
                AS path: 100 I
```

```
                    Localpref: 100
                    Router ID: 10.0.0.0
```

show route inactive-path extensive

The output for the **show route inactive-path extensive** command is identical to that of the **show route inactive-path detail** command. For sample output, see **show route inactive-path detail on page 1961**.

show route inactive-path terse

```
user@host> show route inactive-path terse

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  10.12.100.12/30    O  10          1            >so-0/3/0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  10.0.0.0/8         D   0                     >fxp1.0

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  10.12.80.0/30      B 170        100            >10.12.80.1      100 I

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route inactive-prefix

Syntax
show route inactive-prefix
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show route inactive-prefix
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display inactive route destinations in each routing table.

Options
none—Display all inactive route destination.

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route inactive-prefix on page 1963
show route inactive-prefix detail on page 1963
show route inactive-prefix extensive on page 1964
show route inactive-prefix terse on page 1964

Output Fields
For information about output fields, see the output field tables for the **show route** command, the **show route detail** command, the **show route extensive** command, or the **show route terse** command.

show route inactive-prefix

```
user@host> show route inactive-prefix

inet.0: 14 destinations, 14 routes (13 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

127.0.0.1/32        [Direct/0] 00:04:54
                    > via lo0.0
```

show route inactive-prefix detail

```
user@host> show route inactive-prefix detail

inet.0: 14 destinations, 14 routes (13 active, 0 holddown, 1 hidden)
127.0.0.1/32 (1 entry, 0 announced)
         Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Hidden Martian Int>
                Age: 4:51
                Task: IF
```

```
                    AS path: I00:04:54
                        > via lo0.0
```

show route inactive-prefix extensive

The output for the **show route inactive-prefix extensive** command is identical to that of the **show route inactive-path detail** command. For sample output, see **show route inactive-prefix detail on page 1963**.

show route inactive-prefix terse

```
user@host> show route inactive-prefix terse

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  127.0.0.1/32       D   0                      >lo0.0
```

show route instance

Syntax
show route instance
<brief | detail | summary>
<*instance-name*>
<logical-system (all | *logical-system-name*)>
<operational>

Syntax (J-EX Series Switch)
show route instance
<brief | detail | summary>
<*instance-name*>
<operational>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display routing instance information.

Options
none—(Same as **brief**) Display standard information about all routing instances.

brief | detail | summary—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**. (These options are not available with the **operational** keyword.)

instance-name—(Optional) Display information for a specified routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

operational—(Optional) Display operational routing instances.

Required Privilege Level
view

List of Sample Output
show route instance on page 1966
show route instance detail (Graceful Restart Complete) on page 1967
show route instance detail (Graceful Restart Incomplete) on page 1968
show route instance detail (VPLS Routing Instance) on page 1970
show route instance operational on page 1970
show route instance summary on page 1970

Output Fields
Table 248 on page 1965 lists the output fields for the **show route instance** command. Output fields are listed in the approximate order in which they appear.

Table 248: show route instance Output Fields

Field Name	Field Description	Level of Output
Instance or ***instance-name***	Name of the routing instance.	All levels
Operational Routing Instances	(**operational** keyword only) Names of all operational routing instances.	—
Type	Type of routing instance: **forwarding**, **l2vpn**, **no-forwarding**, **vpls**, or **vrf**.	All levels

Table 248: show route instance Output Fields *(continued)*

Field Name	Field Description	Level of Output
State	State of the routing instance: **active** or **inactive**.	**brief detail** none
Interfaces	Name of interfaces belonging to this routing instance.	**brief detail** none
Restart State	Status of graceful restart for this instance: **Pending** or **Complete**.	**detail**
Path selection timeout	Maximum amount of time, in seconds, remaining until graceful restart is declared complete. The default is 300.	**detail**
Tables	Tables (and number of routes) associated with this routing instance.	none **brief detail**
Route-distinguisher	Unique route distinguisher associated with this routing instance.	**detail**
Vrf-import	VPN routing and forwarding instance import policy name.	**detail**
Vrf-export	VPN routing and forwarding instance export policy name.	**detail**
Vrf-import-target	VPN routing and forwarding instance import target community name.	**detail**
Vrf-export-target	VPN routing and forwarding instance export target community name.	**detail**
Fast-reroute-priority	Fast reroute priority setting for a VPLS routing instance: **high**, **medium**, or **low**. The default is **low**.	**detail**
Restart State	Restart state: • **Pending:*protocol-name***—List of protocols that have not yet completed graceful restart for this routing table. • **Complete**—All protocols have restarted for this routing table.	**detail**
Primary rib	Primary table for this routing instance.	**brief** none **summary**
Active/holddown/hidden	Number of active, hold-down, and hidden routes.	All levels

show route instance

```
user@host> show route instance
Instance             Type
        Primary RIB                                  Active/holddown/hidden
master               forwarding
        inet.0                                       16/0/1
        iso.0                                        1/0/0
        mpls.0                                       0/0/0
        inet6.0                                      2/0/0
        l2circuit.0                                  0/0/0
__juniper_private1__ forwarding
        __juniper_private1__.inet.0                  12/0/0
        __juniper_private1__.inet6.0                 1/0/0
```

show route instance detail (Graceful Restart Complete)

```
user@host> show route instance detail
master:
  Router ID: 10.255.14.176
  Type: forwarding        State: Active
  Restart State: Complete Path selection timeout: 300
  Tables:
    inet.0                : 17 routes (15 active, 0 holddown, 1 hidden)
    Restart Complete
    inet.3                : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    iso.0                 : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
    mpls.0                : 19 routes (19 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l3vpn.0           : 10 routes (10 active, 0 holddown, 0 hidden)
    Restart Complete
    inet6.0               : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l2vpn.0           : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
BGP-INET:
  Router ID: 10.69.103.1
  Type: vrf               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.103
  Route-distinguisher: 10.255.14.176:103
  Vrf-import: [ BGP-INET-import ]
  Vrf-export: [ BGP-INET-export ]
  Tables:
    BGP-INET.inet.0       : 4 routes (4 active, 0 holddown, 0 hidden)
    Restart Complete
BGP-L:
  Router ID: 10.69.104.1
  Type: vrf               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.104
  Route-distinguisher: 10.255.14.176:104
  Vrf-import: [ BGP-L-import ]
  Vrf-export: [ BGP-L-export ]
  Tables:
    BGP-L.inet.0          : 4 routes (4 active, 0 holddown, 0 hidden)
    Restart Complete
    BGP-L.mpls.0          : 3 routes (3 active, 0 holddown, 0 hidden)
    Restart Complete
L2VPN:
  Router ID: 0.0.0.0
  Type: l2vpn             State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.512
  Route-distinguisher: 10.255.14.176:512
  Vrf-import: [ L2VPN-import ]
  Vrf-export: [ L2VPN-export ]
  Tables:
    L2VPN.l2vpn.0         : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
LDP:
  Router ID: 10.69.105.1
  Type: vrf               State: Active
```

```
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.105
  Route-distinguisher: 10.255.14.176:105
  Vrf-import: [ LDP-import ]
  Vrf-export: [ LDP-export ]
  Tables:
    LDP.inet.0              : 5 routes (4 active, 0 holddown, 0 hidden)
    Restart Complete
OSPF:
  Router ID: 10.69.101.1
  Type: vrf               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.101
  Route-distinguisher: 10.255.14.176:101
  Vrf-import: [ OSPF-import ]
  Vrf-export: [ OSPF-export ]
  Vrf-import-target: [ target:11111
  Tables:
    OSPF.inet.0             : 8 routes (7 active, 0 holddown, 0 hidden)
    Restart Complete
RIP:
  Router ID: 10.69.102.1
  Type: vrf               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.102
  Route-distinguisher: 10.255.14.176:102
  Vrf-import: [ RIP-import ]
  Vrf-export: [ RIP-export ]
  Tables:
    RIP.inet.0              : 6 routes (6 active, 0 holddown, 0 hidden)
    Restart Complete
STATIC:
  Router ID: 10.69.100.1
  Type: vrf               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.100
  Route-distinguisher: 10.255.14.176:100
  Vrf-import: [ STATIC-import ]
  Vrf-export: [ STATIC-export ]
  Tables:
    STATIC.inet.0           : 4 routes (4 active, 0 holddown, 0 hidden)
    Restart Complete
```

show route instance detail (Graceful Restart Incomplete)

```
user@host> show route instance detail
master:
  Router ID: 10.255.14.176
  Type: forwarding        State: Active
  Restart State: Pending  Path selection timeout: 300
  Tables:
    inet.0                  : 17 routes (15 active, 1 holddown, 1 hidden)
    Restart Pending: OSPF LDP
    inet.3                  : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: OSPF LDP
    iso.0                   : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
    mpls.0                  : 23 routes (23 active, 0 holddown, 0 hidden)
    Restart Pending: LDP VPN
```

```
    bgp.l3vpn.0            : 10 routes (10 active, 0 holddown, 0 hidden)
    Restart Pending: BGP VPN
    inet6.0                : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l2vpn.0            : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Pending: BGP VPN
BGP-INET:
  Router ID: 10.69.103.1
  Type: vrf               State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.103
  Route-distinguisher: 10.255.14.176:103
  Vrf-import: [ BGP-INET-import ]
  Vrf-export: [ BGP-INET-export ]
  Tables:
    BGP-INET.inet.0        : 6 routes (5 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
BGP-L:
  Router ID: 10.69.104.1
  Type: vrf               State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.104
  Route-distinguisher: 10.255.14.176:104
  Vrf-import: [ BGP-L-import ]
  Vrf-export: [ BGP-L-export ]
  Tables:
    BGP-L.inet.0           : 6 routes (5 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
    BGP-L.mpls.0           : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
L2VPN:
  Router ID: 0.0.0.0
  Type: l2vpn             State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.512
  Route-distinguisher: 10.255.14.176:512
  Vrf-import: [ L2VPN-import ]
  Vrf-export: [ L2VPN-export ]
  Tables:
    L2VPN.l2vpn.0          : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: VPN L2VPN
LDP:
  Router ID: 10.69.105.1
  Type: vrf               State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.105
  Route-distinguisher: 10.255.14.176:105
  Vrf-import: [ LDP-import ]
  Vrf-export: [ LDP-export ]
  Tables:
    LDP.inet.0             : 5 routes (4 active, 1 holddown, 0 hidden)
    Restart Pending: OSPF LDP VPN
OSPF:
  Router ID: 10.69.101.1
  Type: vrf               State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
```

```
    t3-0/0/0.101
  Route-distinguisher: 10.255.14.176:101
  Vrf-import: [ OSPF-import ]
  Vrf-export: [ OSPF-export ]
  Tables:
    OSPF.inet.0            : 8 routes (7 active, 1 holddown, 0 hidden)
    Restart Pending: OSPF VPN
RIP:
  Router ID: 10.69.102.1
  Type: vrf               State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.102
  Route-distinguisher: 10.255.14.176:102
  Vrf-import: [ RIP-import ]
  Vrf-export: [ RIP-export ]
  Tables:
    RIP.inet.0             : 8 routes (6 active, 2 holddown, 0 hidden)
    Restart Pending: RIP VPN
STATIC:
  Router ID: 10.69.100.1
  Type: vrf               State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.100
  Route-distinguisher: 10.255.14.176:100
  Vrf-import: [ STATIC-import ]
  Vrf-export: [ STATIC-export ]
  Tables:
    STATIC.inet.0          : 4 routes (4 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
```

show route instance detail (VPLS Routing Instance)

```
user@host> show route instance detail test-vpls
test-vpls:
  Router ID: 0.0.0.0
  Type: vpls              State: Active
  Interfaces:
    lsi.1048833
    lsi.1048832
    fe-0/1/0.513
  Route-distinguisher: 10.255.37.65:1
  Vrf-import: [ __vrf-import-test-vpls-internal__ ]
  Vrf-export: [ __vrf-export-test-vpls-internal__ ]
  Vrf-import-target: [ target:300:1 ]
  Vrf-export-target: [ target:300:1 ]
  Fast-reroute-priority: high
  Tables:
    test-vpls.l2vpn.0          : 3 routes (3 active, 0 holddown, 0 hidden)
```

show route instance operational

```
user@host> show route instance operational
Operational Routing Instances:

master
default
```

show route instance summary

```
user@host> show route instance summary
Instance           Type           Primary rib            Active/holddown/hidden
master             forwarding
                                  inet.0                 15/0/1
                                  iso.0                  1/0/0
```

```
                                    mpls.0                  35/0/0
                                    l3vpn.0                 0/0/0
                                    inet6.0                 2/0/0
                                    l2vpn.0                 0/0/0
                                    l2circuit.0             0/0/0
BGP-INET            vrf
                                    BGP-INET.inet.0         5/0/0
                                    BGP-INET.iso.0          0/0/0
                                    BGP-INET.inet6.0        0/0/0
BGP-L               vrf
                                    BGP-L.inet.0            5/0/0
                                    BGP-L.iso.0             0/0/0
                                    BGP-L.mpls.0            4/0/0
                                    BGP-L.inet6.0           0/0/0
L2VPN               l2vpn
                                    L2VPN.inet.0            0/0/0
                                    L2VPN.iso.0             0/0/0
                                    L2VPN.inet6.0           0/0/0
                                    L2VPN.l2vpn.0           2/0/0
LDP                 vrf
                                    LDP.inet.0              4/0/0
                                    LDP.iso.0               0/0/0
                                    LDP.mpls.0              0/0/0
                                    LDP.inet6.0             0/0/0
                                    LDP.l2circuit.0         0/0/0
OSPF                vrf
                                    OSPF.inet.0             7/0/0
                                    OSPF.iso.0              0/0/0
                                    OSPF.inet6.0            0/0/0
RIP                 vrf
                                    RIP.inet.0              6/0/0
                                    RIP.iso.0               0/0/0
                                    RIP.inet6.0             0/0/0
STATIC              vrf
                                    STATIC.inet.0           4/0/0
                                    STATIC.iso.0            0/0/0
                                    STATIC.inet6.0          0/0/0
```

show route label

Syntax
show route label *label*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show route label *label*
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display the routes based on a specified Multiprotocol Label Switching (MPLS) label value.

Options
label—Value of the MPLS label.

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route label on page 1972
show route label detail on page 1973
show route label extensive on page 1973
show route label terse on page 1973

Output Fields
For information about output fields, see the output field table for the **show route** command, the **show route detail** command, the **show route extensive** command, or the **show route terse** command.

show route label

```
user@host> show route label 100016

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
100016             *[VPN/170] 03:25:41
                    > to 10.12.80.1 via ge-6/3/2.0, Pop
```

show route label detail

```
user@host> show route label 100016 detail

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
100016 (1 entry, 1 announced)
        *VPN    Preference: 170
                Next-hop reference count: 2
                Source: 10.12.80.1
                Next hop: 10.12.80.1 via ge-6/3/2.0, selected
                Label operation: Pop
                State: <Active Int Ext>
                Local AS:     1
                Age: 3:23:31
                Task: BGP.0.0.0.0+179
                Announcement bits (1): 0-KRT
                AS path: 100 I
                Ref Cnt: 2
```

show route label extensive

The output for the show route label extensive command is identical to that of the **show route label detail** command. For sample output, see **show route label detail on page 1973**.

show route label terse

```
user@host> show route label 100016 terse

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 100016             V 170                       >10.12.80.1
```

show route label-switched-path

Syntax
show route label-switched-path *path-name*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show route label-switched-path *path-name*
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display the routes used in a Multiprotocol Label Switching (MPLS) label-switched path (LSP).

Options
brief | detail | extensive | terse—(Optional) Display the specified level of output.

path-name—LSP tunnel name.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route label-switched-path on page 1974

Output Fields
For information about output fields, see the output field tables for the **show route** command, the **show route detail** command, the **show route extensive** command, or the **show route terse** command.

show route label-switched-path

```
user@host> show route label-switched-path sf-to-ny
inet.0: 29 destinations, 29 routes (29 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.1/32          [MPLS/7] 00:00:06, metric 0
                    > to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
3.3.3.3/32         *[MPLS/7] 00:00:06, metric 0
                    > to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2.2.2.2/32         *[MPLS/7] 00:00:06, metric 0
                    > to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
4.4.4.4/32         *[MPLS/7] 00:00:06, metric 0
                      to 111.222.1.9 via s0-0/0/0, label-switched-path abc
                    > to 111.222.1.9 via s0-0/0/0, label-switched-path xyz
                      to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny
111.222.1.9/32      [MPLS/7] 00:00:06, metric 0
                    > to 111.222.1.9 via s0-0/0/0, label-switched-path sf-to-ny

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

```
mpls.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```

show route martians

Syntax

show route martians
<logical-system (all | *logical-system-name*)>
<table *routing-table-name*>

Syntax (J-EX Series Switch)

show route martians
<table *routing-table-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display the martian (invalid and ignored) entries associated with each routing table.

Options

none—Display standard information about route martians for all routing tables.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

table *routing-table-name*—(Optional) Display only the martian entries associated with a particular routing table.

Required Privilege Level

view

List of Sample Output

show route martians on page 1976

Output Fields

Table 249 on page 1976 lists the output fields for the **show route martians** command. Output fields are listed in the approximate order in which they appear

Table 249: show route martians Output Fields

Field Name	Field Description
table-name	Name of the route table in which the route martians reside.
destination-prefix	Route destination.
match value	Route match parameter.
status	Status of the route: **allowed** or **disallowed**.

show route martians

```
user@host> show route martians

inet.0:
             0.0.0.0/0 exact -- allowed
             0.0.0.0/8 orlonger -- disallowed
             127.0.0.0/8 orlonger -- disallowed
             128.0.0.0/16 orlonger -- disallowed
             191.255.0.0/16 orlonger -- disallowed
             192.0.0.0/24 orlonger -- disallowed
             223.255.255.0/24 orlonger -- disallowed
             240.0.0.0/4 orlonger -- disallowed
```

```
inet.1:
                0.0.0.0/0 exact -- allowed
                0.0.0.0/8 orlonger -- disallowed
                127.0.0.0/8 orlonger -- disallowed
                128.0.0.0/16 orlonger -- disallowed
                191.255.0.0/16 orlonger -- disallowed
                192.0.0.0/24 orlonger -- disallowed
                223.255.255.0/24 orlonger -- disallowed
                240.0.0.0/4 orlonger -- disallowed

....
```

show route next-hop

Syntax

show route next-hop *next-hop*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)

show route next-hop *next-hop*
<brief | detail | extensive | terse>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display the entries in the routing table that are being sent to the specified next-hop address.

Options

brief | detail | extensive | terse—(Optional) Display the specified level of ouput.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

next-hop—Next-hop address.

Required Privilege Level

view

List of Sample Output

show route next-hop on page 1978
show route next-hop detail on page 1979
show route next-hop extensive on page 1980
show route next-hop terse on page 1982

Output Fields

For information about output fields, see the output field tables for the **show route** command, the **show route detail** command, the **show route extensive** command, or the **show route terse** command.

show route next-hop

```
user@host> show route next-hop 192.168.71.254

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16       *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
10.209.0.0/16      *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
172.16.0.0/12      *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
192.168.0.0/16     *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
192.168.102.0/23   *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
207.17.136.0/24    *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
207.17.136.192/32  *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
```

```
private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route next-hop detail

```
user@host> show route next-hop 192.168.71.254 detail

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
Restart Complete
10.10.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

10.209.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

172.16.0.0/12 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

192.168.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
```

```
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

192.168.102.0/23 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

207.17.136.0/24 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

207.17.136.192/32 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route next-hop extensive

```
user@host> show route next-hop 192.168.71.254 extensive

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
```

```
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.209.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

172.16.0.0/12 (1 entry, 1 announced)
TSI:
KRT in-kernel 172.16.0.0/12 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

192.168.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 192.168.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

192.168.102.0/23 (1 entry, 1 announced)
TSI:
KRT in-kernel 192.168.102.0/23 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

207.17.136.0/24 (1 entry, 1 announced)
TSI:
```

```
KRT in-kernel 207.17.136.0/24 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

207.17.136.192/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.192/32 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

green.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

red.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route next-hop terse

```
user@host> show route next-hop 192.168.71.254 terse

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.10.0.0/16       S   5                     >192.168.71.254
* 10.209.0.0/16      S   5                     >192.168.71.254
* 172.16.0.0/12      S   5                     >192.168.71.254
* 192.168.0.0/16     S   5                     >192.168.71.254
* 192.168.102.0/23   S   5                     >192.168.71.254
* 207.17.136.0/24    S   5                     >192.168.71.254
* 207.17.136.192/32  S   5                     >192.168.71.254

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete
```

```
mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route no-community

Syntax

show route no-community
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)

show route no-community
<brief | detail | extensive | terse>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display the route entries in each routing table that are not associated with any community.

Options

none—(Same as **brief**) Display the route entries in each routing table that are not associated with any community.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show route no-community on page 1984
show route no-community detail on page 1985
show route no-community extensive on page 1985
show route no-community terse on page 1986

Output Fields

For information about output fields, see the output field tables for the **show route** command, the **show route detail** command, the **show route extensive** command, or the **show route terse** command.

show route no-community

```
user@host> show route no-community
inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16       *[Static/5] 00:36:27
                    > to 192.168.71.254 via fxp0.0
10.209.0.0/16      *[Static/5] 00:36:27
                    > to 192.168.71.254 via fxp0.0
10.255.71.52/32    *[Direct/0] 00:36:27
                    > via lo0.0
10.255.71.63/32    *[OSPF/10] 00:04:39, metric 1
                    > to 35.1.1.2 via ge-3/1/0.0
10.255.71.64/32    *[OSPF/10] 00:00:08, metric 2
                    > to 35.1.1.2 via ge-3/1/0.0
10.255.71.240/32   *[OSPF/10] 00:05:04, metric 2
                      via so-0/1/2.0
                    > via so-0/3/2.0
10.255.71.241/32   *[OSPF/10] 00:05:14, metric 1
                    > via so-0/1/2.0
10.255.71.242/32   *[OSPF/10] 00:05:19, metric 1
                    > via so-0/3/2.0
```

```
12.1.1.0/24        *[OSPF/10] 00:05:14, metric 2
                    > via so-0/3/2.0
14.1.1.0/24        *[OSPF/10] 00:00:08, metric 3
                    > to 35.1.1.2 via ge-3/1/0.0
                      via so-0/1/2.0
                      via so-0/3/2.0
16.1.1.0/24        *[OSPF/10] 00:05:14, metric 2
                    > via so-0/1/2.0
.....
```

show route no-community detail

```
user@host> show route no-community detail

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 38:08
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.209.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 38:08
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I
....
```

show route no-community extensive

```
user@host> show route no-community extensive

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:03:33
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.209.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:03:33
                Task: RT
                Announcement bits (1): 0-KRT
```

```
                    AS path: I
```

show route no-community terse

```
user@host> show route no-community terse

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.10.0.0/16       S   5                       >192.168.71.254
* 10.209.0.0/16      S   5                       >192.168.71.254
* 10.255.71.52/32    D   0                       >lo0.0
* 10.255.71.63/32    O  10          1            >35.1.1.2
* 10.255.71.64/32    O  10          2            >35.1.1.2
* 10.255.71.240/32   O  10          2             so-0/1/2.0
                                                 >so-0/3/2.0
* 10.255.71.241/32   O  10          1            >so-0/1/2.0
* 10.255.71.242/32   O  10          1            >so-0/3/2.0
* 12.1.1.0/24        O  10          2            >so-0/3/2.0
* 14.1.1.0/24        O  10          3            >35.1.1.2
                                                  so-0/1/2.0
                                                  so-0/3/2.0
* 16.1.1.0/24        O  10          2            >so-0/1/2.0
...
```

show route protocol

Syntax
show route protocol *protocol*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show route protocol *protocol*
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display the route entries in the routing table that were learned from a particular protocol.

Options
protocol—Protocol from which the route was learned:

- **access**—Access route for use by DHCP application
- **access-internal**—Access-internal route for use by DHCP application
- **aggregate**—Locally generated aggregate route
- **atmvpn**—Asynchronous Transfer Mode virtual private network
- **bgp**—Border Gateway Protocol
- **ccc**—Circuit cross-connect
- **direct**—Directly connected route
- **dvmrp**—Distance Vector Multicast Routing Protocol
- **esis**—End System-to-Intermediate System
- **flow**—Locally defined flow-specification route.
- **isis**—Intermediate System-to-Intermediate System
- **ldp**—Label Distribution Protocol
- **l2circuit**—Layer 2 circuit
- **l2vpn**—Layer 2 virtual private network
- **local**—Local address
- **mpls**—Multiprotocol Label Switching
- **msdp**—Multicast Source Discovery Protocol
- **ospf**—Open Shortest Path First versions 2 and 3
- **ospf2**—Open Shortest Path First versions 2 only
- **ospf3**—Open Shortest Path First version 3 only
- **pim**—Protocol Independent Multicast
- **rip**—Routing Information Protocol
- **ripng**—Routing Information Protocol next generation

- **rsvp**—Resource Reservation Protocol
- **rtarget**—Local route target virtual private network
- **static**—Statically defined route
- **tunnel**—Dynamic tunnel
- **vpn**—Virtual private network

NOTE: J-EX Series switches run a subset of these protocols. See the switch CLI for details.

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to brief.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level view

List of Sample Output

show route protocol access on page 1988
show route protocol access-internal extensive on page 1989
show route protocol bgp on page 1989
show route protocol bgp detail on page 1989
show route protocol bgp extensive on page 1989
show route protocol bgp terse on page 1990
show route protocol direct on page 1990
show route protocol l2circuit detail on page 1990
show route protocol l2vpn extensive on page 1991
show route protocol ldp on page 1992
show route protocol ldp extensive on page 1992
show route protocol ospf (Layer 3 VPN) on page 1993
show route protocol ospf detail on page 1994
show route protocol rip on page 1994
show route protocol rip detail on page 1994
show route protocol ripng table inet6 on page 1994

Output Fields For information about output fields, see the output field tables for the **show route** command, the **show route detail** command, the **show route extensive** command, or the **show route terse** command.

show route protocol access

```
user@host> show route protocol access

inet.0: 30380 destinations, 30382 routes (30379 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

13.160.0.3/32      *[Access/13] 00:00:09
                    > to 13.160.0.2 via fe-0/0/0.0
13.160.0.4/32      *[Access/13] 00:00:09
                    > to 13.160.0.2 via fe-0/0/0.0
```

```
13.160.0.5/32      *[Access/13] 00:00:09
                    > to 13.160.0.2 via fe-0/0/0.0
```

show route protocol access-internal extensive

```
user@host> show route protocol access-internal 13.160.0.19 extensive

inet.0: 100020 destinations, 100022 routes (100019 active, 0 holddown, 1 hidden)
13.160.0.19/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 13.160.0.19/32 -> {13.160.0.2}
        *Access-internal Preference: 12
                Next-hop reference count: 200000
                Next hop: 13.160.0.2 via fe-0/0/0.0, selected
                State: <Active Int>
       Age: 36
                Task: RPD Unix Domain Server./var/run/rpd_serv.local
                Announcement bits (1): 0-KRT
                AS path: I
```

show route protocol bgp

```
user@host> show route protocol bgp 192.168.64.0/21

inet.0: 24 destinations, 32 routes (23 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.64.0/21     [BGP/170] 00:04:33, localpref 100
                      AS path: 10023 21 I
                    > to 100.1.3.2 via ge-5/0/3.0, Push 100080
```

show route protocol bgp detail

```
show route protocol bgp 66.117.63.0/24 exact detail
inet.0: 227318 destinations, 227319 routes (227305 active, 0 holddown, 13 hidden)
66.117.63.0/24 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Indirect
                Next-hop reference count: 681816
                Source: 207.17.136.192
                Next hop type: Router, Next hop index: 324
                Next hop: 192.168.167.254 via fxp0.0, selected
                Protocol next hop: 207.17.136.29
                Indirect next hop: 8c7b09c 342
                State: <Active Int Ext
                Local AS:   200 Peer AS: 10458
                Age: 20:31:24   Metric2: 0
                Task: BGP_10458_10458.207.17.136.192+179
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: AS2 PA[6]: 14203 2914 3356 29748 33437 AS_TRANS
                AS path: AS4 PA[2]: 33437 393219
                AS path: Merged[6]: 14203 2914 3356 29748 33437 393219 I
                Communities: 2914:420
                Localpref: 100
                Router ID: 207.17.136.192
```

show route protocol bgp extensive

```
user@host> show route protocol bgp 192.168.64.0/21 extensive

inet.0: 24 destinations, 32 routes (23 active, 0 holddown, 1 hidden)
192.168.64.0/21 (2 entries, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 86f50a8
         BGP    Preference: 170/-101
                Next-hop reference count: 3
                Source: 100.1.3.2
```

```
                    Next hop: 100.1.3.2 via ge-5/0/3.0, selected
                    Label operation: Push 100080
                    State: <Ext>
                    Inactive reason: Route Preference
                    Local AS:    21 Peer AS: 10023
                    Age: 4:43
                    Task: BGP_10023.100.1.3.2+4282
                    AS path: 10023 21 I
                    Route Label: 100080
                    Localpref: 100
                    Router ID: 100.1.3.2
```

show route protocol bgp terse

```
user@host> show route protocol bgp 192.168.64.0/21 terse

inet.0: 24 destinations, 32 routes (23 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  192.168.64.0/21    B 170        100            >100.1.3.2       10023 21 I
```

show route protocol direct

```
user@host> show route protocol direct

inet.0: 35 destinations, 35 routes (34 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

127.0.0.1/32        [Direct/0] 14:36:24
                     > via lo0.0
111.222.5.0/24     *[Direct/0] 14:36:24
                     > via fxp0.0
111.222.8.16/28    *[Direct/0] 14:36:24
                     > via at-5/3/0.0
111.222.8.100/30   *[Direct/0] 14:36:24
                     > via at-5/3/0.129
111.222.8.104/30   *[Direct/0] 14:36:24
                     > via at-5/3/0.128
111.222.8.161/32   *[Direct/0] 14:36:24
                     > via t3-5/2/0.0
111.222.8.163/32   *[Direct/0] 14:36:24
                     > via t3-5/2/1.0
...

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

47.0005.80ff.f800.0000.0108.0001.1921.6800.5081.00/160
                   *[Direct/0] 14:36:24
                     > via lo0.0
```

show route protocol l2circuit detail

```
user@host> show route protocol l2circuit detail

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
100000 (1 entry, 1 announced)
        *L2CKT  Preference: 7
                Next hop: via ge-2/0/0.0, selected
                Label operation: Pop       Offset: 4
                State: <Active Int>
                Local AS:    99
                Age: 9:52
                Task: Common L2 VC
                Announcement bits (1): 0-KRT
```

```
                AS path: I

ge-2/0/0.0 (1 entry, 1 announced)
        *L2CKT  Preference: 7
                Next hop: via so-1/1/2.0 weight 1, selected
                Label-switched-path my-lsp
                Label operation: Push 100000, Push 100000(top)[0] Offset: -4
                Protocol next hop: 10.245.255.63
                Push 100000 Offset: -4
                 Indirect next hop: 86af0c0 298
                State: <Active Int>
                Local AS:    99
                Age: 9:52
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 1-Common L2 VC
                AS path: I

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
        *L2CKT  Preference: 7
                Next hop: via so-1/1/2.0 weight 1, selected
                Label-switched-path my-lsp
                Label operation: Push 100000[0]
                Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
                State: <Active Int>
                Local AS:    99
                Age: 10:21
                Task: l2 circuit
                Announcement bits (1): 0-LDP
                AS path: I
                VC Label 100000, MTU 1500, VLAN ID 512
```

show route protocol l2vpn extensive

```
user@host> show route protocol l2vpn extensive

inet.0: 14 destinations, 15 routes (13 active, 0 holddown, 1 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
800001 (1 entry, 1 announced)
TSI:
KRT in-kernel 800001 /36 -> {so-0/0/0.0}
        *L2VPN  Preference: 7
                Next hop: via so-0/0/0.0 weight 49087 balance 97%, selected
                Label operation: Pop       Offset: 4
                State: <Active Int>
                Local AS:    69
                Age: 7:48
                Task: Common L2 VC
                Announcement bits (1): 0-KRT
                AS path: I

so-0/0/0.0 (1 entry, 1 announced)
TSI:
KRT in-kernel so-0/0/0.0.0      /16 -> {indirect(288)}
        *L2VPN  Preference: 7
                Next hop: via so-0/0/1.0, selected
                Label operation: Push 800000 Offset: -4
```

```
                    Protocol next hop: 10.255.14.220
                    Push 800000 Offset: -4
                     Indirect next hop: 85142a0 288
                    State: <Active Int>
                    Local AS:    69
                    Age: 7:48
                    Task: Common L2 VC
                    Announcement bits (2): 0-KRT 1-Common L2 VC
                    AS path: I
                    Communities: target:69:1 Layer2-info: encaps:PPP,
                    control flags:2, mtu: 0
```

show route protocol ldp

```
user@host> show route protocol ldp
inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.16.1/32    *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0, Push 100000
192.168.17.1/32    *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0

private1__.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100064             *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0, Pop
100064(S=0)        *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0, Pop
100080             *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0, Swap 100000
```

show route protocol ldp extensive

```
user@host> show route protocol ldp extensive
192.168.16.1/32 (1 entry, 1 announced)
        State: <FlashAll>
        *LDP    Preference: 9
                Next-hop reference count: 3
                Next hop: via t1-4/0/0.0, selected
                Label operation: Push 100000
                State: <Active Int>
                Local AS: 65500
                Age: 1d 23:03:58        Metric: 1
                Task: LDP
                Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
                AS path: I

192.168.17.1/32 (1 entry, 1 announced)
        State: <FlashAll>
        *LDP    Preference: 9
                Next-hop reference count: 3
                Next hop: via t1-4/0/0.0, selected
                State: <Active Int>
                Local AS: 65500
                Age: 1d 23:03:58        Metric: 1
                Task: LDP
                Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
                AS path: I
```

```
private1__.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)

100064 (1 entry, 1 announced)
TSI:
KRT in-kernel 100064 /36 -> {t1-4/0/0.0}
        *LDP    Preference: 9
                Next-hop reference count: 2
                Next hop: via t1-4/0/0.0, selected
                State: <Active Int>
                Local AS: 65500
                Age: 1d 23:03:58        Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
                Prefixes bound to route: 192.168.17.1/32

100064(S=0) (1 entry, 1 announced)
TSI:
KRT in-kernel 100064 /40 -> {t1-4/0/0.0}
        *LDP    Preference: 9
                Next-hop reference count: 2
                Next hop: via t1-4/0/0.0, selected
                Label operation: Pop
                State: <Active Int>
                Local AS: 65500
                Age: 1d 23:03:58        Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I

100080 (1 entry, 1 announced)
TSI:
KRT in-kernel 100080 /36 -> {t1-4/0/0.0}
        *LDP    Preference: 9
                Next-hop reference count: 2
                Next hop: via t1-4/0/0.0, selected
                Label operation: Swap 100000
                State: <Active Int>
                Local AS: 65500
                Age: 1d 23:03:58        Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
                Prefixes bound to route: 192.168.16.1/32
```

show route protocol ospf (Layer 3 VPN)

```
user@host> show route protocol ospf
inet.0: 40 destinations, 40 routes (39 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.4/30       *[OSPF/10] 00:05:18, metric 4
                    > via t3-3/2/0.0
10.39.1.8/30        [OSPF/10] 00:05:18, metric 2
                    > via t3-3/2/0.0
10.255.14.171/32   *[OSPF/10] 00:05:18, metric 4
                    > via t3-3/2/0.0
10.255.14.179/32   *[OSPF/10] 00:05:18, metric 2
                    > via t3-3/2/0.0
224.0.0.5/32       *[OSPF/10] 20:25:55, metric 1
```

```
VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30       [OSPF/10] 00:05:43, metric 1
                    > via so-0/2/2.0
10.255.14.173/32   *[OSPF/10] 00:05:43, metric 1
                    > via so-0/2/2.0
224.0.0.5/32       *[OSPF/10] 20:26:20, metric 1
```

show route protocol ospf detail

```
user@host> show route protocol ospf detail
VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30 (2 entries, 0 announced)
         OSPF   Preference:  10
                Nexthop: via so-0/2/2.0, selected
                State: <Int>
                Inactive reason: Route Preference
                Age: 6:25       Metric: 1
                Area: 0.0.0.0
                Task: VPN-AB-OSPF
                AS path: I
                Communities: Route-Type:0.0.0.0:1:0

...
```

show route protocol rip

```
user@host> show route protocol rip
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32   *[RIP/100] 20:24:34, metric 2
                    > to 10.39.1.22 via t3-0/2/2.0
224.0.0.9/32       *[RIP/100] 00:03:59, metric 1
```

show route protocol rip detail

```
user@host> show route protocol rip detail
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32 (1 entry, 1 announced)
        *RIP    Preference: 100
                Nexthop: 10.39.1.22 via t3-0/2/2.0, selected
                State: <Active Int>
                Age: 20:25:02   Metric: 2
                Task: VPN-AB-RIPv2
                Announcement bits (2): 0-KRT 2-BGP.0.0.0.0+179
                AS path: I
                Route learned from 10.39.1.22 expires in 96 seconds
```

show route protocol ripng table inet6

```
user@host> show route protocol ripng table inet6
inet6.0: 4215 destinations, 4215 routes (4214 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

1111::1/128        *[RIPng/100] 02:13:33, metric 2
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::2/128        *[RIPng/100] 02:13:33, metric 2
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
```

```
1111::3/128        *[RIPng/100] 02:13:33, metric 2
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::4/128        *[RIPng/100] 02:13:33, metric 2
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::5/128        *[RIPng/100] 02:13:33, metric 2
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::6/128        *[RIPng/100] 02:13:33, metric 2
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
```

show route range

Syntax

show route range
<brief | detail | extensive | terse>
<*destination-prefix*>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)

show route range
<brief | detail | extensive | terse>
<*destination-prefix*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display routing table entries using a prefix range.

Options

none—Display standard information about all routing table entries using a prefix range.

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

destination-prefix—(Optional) Destination and prefix mask for the range.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show route range on page 1996
show route range destination-prefix on page 1997
show route range detail on page 1997
show route range extensive on page 1998
show route range terse on page 1999

Output Fields

For information about output fields, see the output field tables for the **show route** command, the **show route detail** command, the **show route extensive** command, or the **show route terse** command.

show route range

```
user@host> show route range

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16       *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
10.209.0.0/16      *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
10.255.71.14/32    *[Direct/0] 00:30:01
                    > via lo0.0
172.16.0.0/12      *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
192.168.0.0/16     *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
192.168.64.0/21    *[Direct/0] 00:30:01
```

```
                    > via fxp0.0
192.168.71.14/32   *[Local/0] 00:30:01
                      Local via fxp0.0
192.168.102.0/23   *[Static/5] 00:30:01
                    > to 192.168.71.254 via fxp0.0
...
```

show route range destination-prefix

```
user@host> show route range 192.168.0.0

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.0.0/16     *[Static/5] 00:31:14
                    > to 192.168.71.254 via fxp0.0
192.168.64.0/21    *[Direct/0] 00:31:14
                    > via fxp0.0
192.168.71.14/32   *[Local/0] 00:31:14
                      Local via fxp0.0
192.168.102.0/23   *[Static/5] 00:31:14
                    > to 192.168.71.254 via fxp0.0
```

show route range detail

```
user@host> show route range detail

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 30:05
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.209.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 30:05
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.255.71.14/32 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Age: 30:05
                Task: IF
                AS path: I

172.16.0.0/12 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 30:05
                Task: RT
```

```
                Announcement bits (1): 0-KRT
                AS path: I

...
```

show route range extensive

```
user@host> show route range extensive

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 30:17
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.209.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 30:17
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.255.71.14/32 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Age: 30:17
                Task: IF
                AS path: I

172.16.0.0/12 (1 entry, 1 announced)
TSI:
KRT in-kernel 172.16.0.0/12 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Age: 30:17
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

...
```

show route range terse

```
user@host> show route range terse

inet.0: 11 destinations, 11 routes (10 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.10.0.0/16       S   5                     >192.168.71.254
* 10.209.0.0/16      S   5                     >192.168.71.254
* 10.255.71.14/32    D   0                     >lo0.0
* 172.16.0.0/12      S   5                     >192.168.71.254
* 192.168.0.0/16     S   5                     >192.168.71.254
* 192.168.64.0/21    D   0                     >fxp0.0
* 192.168.71.14/32   L   0                      Local
* 192.168.102.0/23   S   5                     >192.168.71.254
* 207.17.136.0/24    S   5                     >192.168.71.254
* 207.17.136.192/32  S   5                     >192.168.71.254

__juniper_private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.0.0.0/8         D   0                     >fxp2.0
                     D   0                     >fxp1.0
* 10.0.0.4/32        L   0                      Local

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  47.0005.80ff.f800.0000.0108.0001.0102.5507.1014/152
*                    D   0                     >lo0.0

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  abcd::10:255:71:14/128
*                    D   0                     >lo0.0
  fe80::280:42ff:fe11:226f/128
*                    D   0                     >lo0.0

__juniper_private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  fe80::280:42ff:fe11:226f/128
*                    D   0                     >lo0.16385
```

show route receive-protocol

Syntax show route receive-protocol *protocol neighbor-address*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)

Syntax (J-EX Series Switch) show route receive-protocol *protocol neighbor-address*
<brief | detail | extensive | terse>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display the routing information as it was received through a particular neighbor using a particular dynamic routing protocol.

Options brief | detail | extensive | terse—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

protocol neighbor-address—Protocol transmitting the route (**bgp**, **dvmrp**, **msdp**, **pim**, **rip**, or **ripng**) and address of the neighboring router from which the route entry was received.

Additional Information The output displays the selected routes and the attributes with which they were received, but does not show the effects of import policy on the routing attributes.

Required Privilege Level view

List of Sample Output
show route receive-protocol bgp on page 2002
show route receive-protocol bgp extensive on page 2002
show route receive-protocol bgp extensive on page 2003
show route receive-protocol bgp detail (Layer 2 VPN) on page 2004
show route receive-protocol bgp extensive (Layer 2 VPN) on page 2004
show route receive-protocol bgp (Layer 3 VPN) on page 2005
show route receive-protocol bgp detail (Layer 3 VPN) on page 2005
show route receive-protocol bgp extensive (Layer 3 VPN) on page 2006

Output Fields Table 250 on page 2000 describes the output fields for the **show route receive-protocol** command. Output fields are listed in the approximate order in which they appear.

Table 250: show route receive-protocol Output Fields

Field Name	Field Description	Level of Output
routing-table-name	Name of the routing table—for example, **inet.0**.	All levels
number **destinations**	Number of destinations for which there are routes in the routing table.	All levels

Table 250: show route receive-protocol Output Fields *(continued)*

Field Name	Field Description	Level of Output
number **routes**	Number of routes in the routing table and total number of routes in the following states: • **active** • **holddown** (routes in that are pending state before being declared inactive) • **hidden** (the routes are not used because of a routing policy)	All levels
Prefix	Destination prefix.	none **brief**
MED	Multiple exit discriminator value included in the route.	none **brief**
destination-prefix **(entry, announced)**	Destination prefix. The **entry** value is the number of routes for this destination, and the **announced** value is the number of routes being announced for this destination.	**detail extensive**
Route Distinguisher	64-bit prefix added to IP subnets to make them unique.	**detail extensive**
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.	**detail extensive**
VPN Label	Virtual private network (VPN) label. Packets are sent between CE and PE routing devices by advertising VPN labels. VPN labels transit over either a Resource Reservation Protocol (RSVP) or a Label Distribution Protocol (LDP) label-switched path (LSP) tunnel.	**detail extensive**
Next hop	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.	All levels
Localpref or **Lclpref**	Local preference value included in the route.	All levels

Table 250: show route receive-protocol Output Fields *(continued)*

Field Name	Field Description	Level of Output
AS path	Autonomous system (AS) path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used the AS-path merge process, as defined in RFC 4893. • **[ ]**—If more than one AS number is configured on the router, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set.	All levels
Cluster list	(For route reflected output only) Cluster ID sent by the route reflector.	**detail extensive**
Originator ID	(For route reflected output only) Address of routing device that originally sent the route to the route reflector.	**detail extensive**
Communities	Community path attribute for the route. See the Output Field table in the **show route detail** command for all possible values for this field.	**detail extensive**
Attrset AS	Number, local preference, and path of the AS that originated the route. These values are stored in the **Attrset** attribute at the originating routing device.	**detail extensive**
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).	**detail extensive**
control flags	Control flags: **none** or **Site Down**.	**detail extensive**
mtu	Maximum transmission unit (MTU) of the Layer 2 circuit.	**detail extensive**

show route receive-protocol bgp

```
user@host> show route receive-protocol bgp 10.255.245.215

inet.0: 28 destinations, 33 routes (27 active, 0 holddown, 1 hidden)
Prefix                  Next hop              MED     Lclpref    AS path
10.22.1.0/24            10.255.245.215        0       100        I
10.22.2.0/24            10.255.245.215        0       100        I
```

show route receive-protocol bgp extensive

```
user@host> show route receive-protocol bgp 10.255.245.63 extensive
inet.0: 244 destinations, 244 routes (243 active, 0 holddown, 1 hidden)
Prefix             Next hop               MED    Lclpref AS path
1.1.1.0/24 (1 entry, 1 announced)
```

```
     Next hop: 10.0.50.3
     Localpref: 100
     AS path: I <Orginator>
     Cluster list:  10.2.3.1
     Originator ID: 10.255.245.45
165.3.0.0/16 (1 entry, 1 announced)
     Next hop: 111.222.5.254
     Localpref: 100
     AS path: I <Originator>
     Cluster list:  10.2.3.1
     Originator ID: 10.255.245.68
165.4.0.0/16 (1 entry, 1 announced)
     Next hop: 111.222.5.254
     Localpref: 100
     AS path: I <Originator>
     Cluster list:  10.2.3.1
     Originator ID: 10.255.245.45
195.1.2.0/24 (1 entry, 1 announced)
     Next hop: 111.222.5.254
     Localpref: 100
     AS path: I <Originator>
     Cluster list:  10.2.3.1
     Originator ID: 10.255.245.68
inet.2: 63 destinations, 63 routes (63 active, 0 holddown, 0 hidden)
Prefix             Next hop                MED    Lclpref AS path
inet.3: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix             Next hop                MED    Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix             Next hop                MED    Lclpref AS path
mpls.0: 48 destinations, 48 routes (48 active, 0 holddown, 0 hidden)
```

show route receive-protocol bgp extensive

```
user@host> show route receive-protocol bgp 207.17.136.192 table inet.0 66.117.68.0/24 extensive
inet.0: 227315 destinations, 227316 routes (227302 active, 0 holddown, 13 hidden)
* 66.117.63.0/24 (1 entry, 1 announced)
     Nexthop: 207.17.136.29
     Localpref: 100
     AS path: AS2 PA[6]: 14203 2914 3356 29748 33437 AS_TRANS
     AS path: AS4 PA[2]: 33437 393219
     AS path: Merged[6]: 14203 2914 3356 29748 33437 393219 I
     Communities: 2914:420
```

show route receive-protocol bgp detail (Layer 2 VPN)

```
user@host> show route receive-protocol bgp 10.255.14.171 detail
inet.0: 68 destinations, 68 routes (67 active, 0 holddown, 1 hidden)
Prefix             Nexthop               MED    Lclpref AS path
inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
frame-vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0
hidden)
Prefix             Nexthop               MED    Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 1 announced)
     Route Distinguisher: 10.255.245.35:1
     Label-base : 800000, range : 4, status-vector : 0x0
     Nexthop: 10.255.245.35
     Localpref: 100
     AS path: I
     Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
     control flags: 0, mtu: 0
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 0 announced)
     Route Distinguisher: 10.255.245.35:1
     Label-base : 800000, range : 4, status-vector : 0x0
     Nexthop: 10.255.245.35
     Localpref: 100
     AS path: I
     Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
     control flags:0, mtu: 0
```

show route receive-protocol bgp extensive (Layer 2 VPN)

```
user@host> show route receive-protocol bgp 10.255.14.171 extensive
inet.0: 68 destinations, 68 routes (67 active, 0 holddown, 1 hidden)
Prefix             Nexthop               MED    Lclpref AS path
inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
frame-vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 1 announced)
     Route Distinguisher: 10.255.245.35:1
     Label-base : 800000, range : 4, status-vector : 0x0
     Nexthop: 10.255.245.35
     Localpref: 100
     AS path: I
     Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
     control flags:0, mtu: 0
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 0 announced)
     Route Distinguisher: 10.255.245.35:1
     Label-base : 800000, range : 4, status-vector : 0x0
     Nexthop: 10.255.245.35
     Localpref: 100
     AS path: I
     Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
     control flags:0, mtu: 0
```

show route receive-protocol bgp (Layer 3 VPN)

```
user@host> show route receive-protocol bgp 10.255.14.171
      inet.0: 33 destinations, 33 routes (32 active, 0 holddown, 1 hidden)
      Prefix             Nexthop               MED    Lclpref AS path
      inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
      Prefix             Nexthop               MED    Lclpref AS path
      VPN-A.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
      Prefix             Nexthop               MED    Lclpref AS path
      10.255.14.175/32   10.255.14.171                    100 2 I
      10.255.14.179/32   10.255.14.171          2         100 I
      VPN-B.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
      Prefix             Nexthop               MED    Lclpref AS path
      10.255.14.175/32   10.255.14.171                    100 2 I
      10.255.14.177/32   10.255.14.171                    100 I
      iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
      Prefix             Nexthop               MED    Lclpref AS path
      mpls.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
      Prefix             Nexthop               MED    Lclpref AS path
      bgp.l3vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
      Prefix             Nexthop               MED    Lclpref AS path
      10.255.14.171:300:10.255.14.177/32
                         10.255.14.171                    100 I
      10.255.14.171:100:10.255.14.179/32
                         10.255.14.171          2         100 I
      10.255.14.171:200:10.255.14.175/32
                         10.255.14.171                    100 2 I
```

show route receive-protocol bgp detail (Layer 3 VPN)

```
user@host> show route receive-protocol bgp 10.255.14.174 detail
inet.0: 16 destinations, 17 routes (15 active, 0 holddown, 1 hidden)
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
vpna.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
* 10.49.0.0/30 (1 entry, 1 announced)
     Route Distinguisher: 10.255.14.176:2
     VPN Label: 101264
     Nexthop: 10.255.14.174
     Localpref: 100
     AS path: I
     Communities: target:200:100
     AttrSet AS: 100
         Localpref: 100
         AS path: I
* 10.255.14.172/32 (1 entry, 1 announced)
     Route Distinguisher: 10.255.14.176:2
     VPN Label: 101280
     Nexthop: 10.255.14.174
     Localpref: 100
     AS path: I
     Communities: target:200:100
     AttrSet AS: 100
         Localpref: 100
         AS path: I
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
* 10.255.14.174:2:10.49.0.0/30 (1 entry, 0 announced)
     Route Distinguisher: 10.255.14.174:2
     VPN Label: 101264
     Nexthop: 10.255.14.174
     Localpref: 100
     AS path: I
     Communities: target:200:100
```

```
     AttrSet AS: 100
         Localpref: 100
         AS path: I
* 10.255.14.174:2:10.255.14.172/32 (1 entry, 0 announced)
     Route Distinguisher: 10.255.14.174:2
     VPN Label: 101280
     Nexthop: 10.255.14.174
     Localpref: 100
     AS path: I
     Communities: target:200:100
     AttrSet AS: 100
         Localpref: 100
         AS path: I
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
```

show route receive-protocol bgp extensive (Layer 3 VPN)

```
user@host> show route receive-protocol bgp 10.255.245.63 extensive
inet.0: 244 destinations, 244 routes (243 active, 0 holddown, 1 hidden)
     Prefix             Nexthop               MED    Lclpref AS path
     1.1.1.0/24 (1 entry, 1 announced)
          Nexthop: 10.0.50.3
          Localpref: 100
          AS path: I <Originator>
          Cluster list:  10.2.3.1
          Originator ID: 10.255.245.45
     165.3.0.0/16 (1 entry, 1 announced)
          Nexthop: 111.222.5.254
          Localpref: 100
          AS path: I <Originator>
          Cluster list:  10.2.3.1
          Originator ID: 10.255.245.68
     165.4.0.0/16 (1 entry, 1 announced)
          Nexthop: 111.222.5.254
          Localpref: 100
          AS path: I <Originator>
          Cluster list:  10.2.3.1
          Originator ID: 10.255.245.45
     195.1.2.0/24 (1 entry, 1 announced)
          Nexthop: 111.222.5.254
          Localpref: 100
          AS path: I <Originator>
          Cluster list:  10.2.3.1
          Originator ID: 10.255.245.68
     inet.2: 63 destinations, 63 routes (63 active, 0 holddown, 0 hidden)
     Prefix             Nexthop               MED    Lclpref AS path
     inet.3: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
     Prefix             Nexthop               MED    Lclpref AS path
     iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
     Prefix             Nexthop               MED    Lclpref AS path
     mpls.0: 48 destinations, 48 routes (48 active, 0 holddown, 0 hidden)
```

show route resolution

Syntax

show route resolution
<brief | detail | extensive | summary>
<index *index*>
<logical-system (all | *logical-system-name*)>
<*prefix*>
<table *routing-table-name*>
<unresolved>

Syntax (J-EX Series Switch)

show route resolution
<brief | detail | extensive | summary>
<index *index*>
<*prefix*>
<table *routing-table-name*>
<unresolved>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display the entries in the next-hop resolution database. This database provides for recursive resolution of next hops through other prefixes in the routing table.

Options

none—Display standard information about all entries in the next-hop resolution database.

brief | detail | extensive | summary—(Optional) Display the specified level of output.

index *index*—(Optional) Show the index of the resolution tree.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

prefix network/destination-prefix—(Optional) Display database entries for the specified address.

table *routing-table-name*—(Optional) Display information about a particular routing table (for example, **inet.0**) where policy-based export is currently enabled. (For information about the different types of routing tables, see the *Junos OS Routing Protocols Configuration Guide*.)

unresolved—(Optional) Display routes that could not be resolved.

Required Privilege Level

view

List of Sample Output

show route resolution detail on page 2008
show route resolution summary on page 2009
show route resolution unresolved on page 2009

Output Fields Table 251 on page 2008 describes the output fields for the **show route resolution** command. Output fields are listed in the approximate order in which they appear.

Table 251: show route resolution Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table whose prefixes are resolved using the entries in the route resolution database. For routing table groups, this is the name of the primary routing table whose prefixes are resolved using the entries in the route resolution database.
Tree index	Tree index identifier.
Nodes	Number of nodes in the tree.
Reference count	Number of references made to the next hop.
Contributing routing tables	Routing tables used for next-hop resolution.
Originating RIB	Name of the routing table whose active route was used to determine the forwarding next-hop entry in the resolution database. For example, in the case of **inet.0** resolving via **inet.0** and **inet.3**, this field indicates which routing table, **inet.0** or **inet.3**, provided the best path for a particular prefix.
Metric	Metric associated with the forwarding next hop.
Node path count	Number of nodes in the path.
Forwarding next hops	Number of forwarding next hops. The forwarding next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.

show route resolution detail

```
user@host> show route resolution detail
Tree Index: 1, Nodes 0, Reference Count 1
Contributing routing tables: inet.3
Tree Index: 2, Nodes 23, Reference Count 1
Contributing routing tables: inet.0 inet.3
10.10.0.0/16 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
10.31.1.0/30 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
10.31.1.1/32 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 0
10.31.1.4/30 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
10.31.1.5/32 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 0
10.31.2.0/30 Originating RIB: inet.0
```

```
  Metric: 2 Node path count: 1
  Forwarding nexthops: 2
10.31.11.0/24 Originating RIB: inet.0
  Node path count: 1
  Forwarding nexthops: 1
```

show route resolution summary

```
user@host> show route resolution summary
Tree Index: 1, Nodes 24, Reference Count 1
 Contributing routing tables: :voice.inet.0 :voice.inet.3
 Tree Index: 2, Nodes 2, Reference Count 1
 Contributing routing tables: inet.3
 Tree Index: 3, Nodes 43, Reference Count 1
 Contributing routing tables: inet.0 inet.3
```

show route resolution unresolved

```
user@host> show route resolution unresolved
Tree Index 1
vt-3/2/0.32769.0      /16
        Protocol Nexthop: 10.255.71.238 Push 800000
        Indirect nexthop: 0 -
vt-3/2/0.32772.0      /16
        Protocol Nexthop: 10.255.70.103 Push 800008
        Indirect nexthop: 0 -
Tree Index 2
```

show route snooping

Syntax

show route snooping
<brief | detail | extensive | terse>
<all>
<best *address/prefix*>
<exact *address*>
<range *prefix-range*>
<summary>
<table *table-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display the entries in the routing table that were learned from snooping.

Options

none—Display the entries in the routing table that were learned from snooping.

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

all—(Optional) Display all entries, including hidden entries.

best *address/prefix*—(Optional) Display the longest match for the provided address and optional prefix.

exact *address/prefix*—(Optional) Display exact matches for the provided address and optional prefix.

range *prefix-range*—(Optional) Display information for the provided address range.

summary—(Optional) Display route snooping summary statisitics.

table *table-name*—(Optional) Display information for the named table.

Required Privilege Level

view

List of Sample Output

show route snooping detail on page 2010

Output Fields

For information about output fields, see the output field tables for the **show route** command, the **show route detail** command, the **show route extensive** command, or the **show route terse** command.

show route snooping detail

```
user@host> show route snooping detail
__+domainAll__.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
 224.0.0.2/32 (1 entry, 1 announced)
        *IGMP   Preference: 0
                Next hop type: MultiRecv
                Next-hop reference count: 4
                State: <Active NoReadvrt Int>
                Age: 2:24
                Task: IGMP
                Announcement bits (1): 0-KRT
                AS path: I
```

```
224.0.0.22/32 (1 entry, 1 announced)
        *IGMP   Preference: 0
                Next hop type: MultiRecv
                Next-hop reference count: 4
                State: <Active NoReadvrt Int>
                Age: 2:24
                Task: IGMP
                Announcement bits (1): 0-KRT
                AS path: I

__+domainAll__.inet.1: 36 destinations, 36 routes (36 active, 0 holddown, 0 hidden)

224.0.0.0.0.0.0.0/24 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4), Next hop index: 1048584
                Next-hop reference count: 4
                State: <Active Int>
                Age: 2:24
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

225.0.0.2.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
                State: <Active Int>
                Age: 2:13
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

225.0.0.3.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
                State: <Active Int>
                Age: 2:15
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

225.0.0.4.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
                State: <Active Int>
                Age: 2:17
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

225.0.0.5.11.11.11.100.3.9.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
                State: <Active Int>
                Age: 1:58
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I
```

```
225.0.0.6,11.11.11.11,100.3.9.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
                State: <Active Int>
                Age: 2:14
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

225.0.0.7,11.11.11.11,100.3.9.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
                State: <Active Int>
                Age: 2:12
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

225.0.0.9,11.11.11.11,100.3.9.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
                State: <Active Int>
                Age: 2:13
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

225.0.0.10,11.11.11.11,100.3.9.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
                State: <Active Int>
                Age: 2:15
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

226.0.0.1,11.11.11.11,100.3.10.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
                State: <Active Int>
                Age: 2:09
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

226.0.0.2,11.11.11.11,100.3.10.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
                State: <Active Int>
                Age: 8
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I
```

```
226.0.0.4.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
                State: <Active Int>
                Age: 2:10
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

226.0.0.8.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
                State: <Active Int>
                Age: 2:12
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

226.0.0.10.11.11.11.100.3.10.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
                State: <Active Int>
                Age: 1:56
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

227.0.0.1.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
                State: <Active Int>
                Age: 2:10
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

227.0.0.2.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
                State: <Active Int>
                Age: 2:13
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

227.0.0.3.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
                State: <Active Int>
                Age: 2:16
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

227.0.0.4.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
```

```
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
                State: <Active Int>
                Age: 2:15
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

227.0.0.5.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
                State: <Active Int>
                Age: 1:57
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

227.0.0.7.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
                State: <Active Int>
                Age: 1:57
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

227.0.0.8.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
                State: <Active Int>
                Age: 2:10
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

227.0.0.10.11.11.11.100.3.11.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
                State: <Active Int>
                Age: 2:15
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

228.0.0.1.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
                State: <Active Int>
                Age: 2:09
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

228.0.0.2.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
```

```
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
                State: <Active Int>
                Age: 2:18
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

228.0.0.7.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
                State: <Active Int>
                Age: 2:11
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

228.0.0.8.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
                State: <Active Int>
                Age: 2:17
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

228.0.0.9.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
                State: <Active Int>
                Age: 8
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

228.0.0.10.11.11.11.100.3.12.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
                State: <Active Int>
                Age: 2:12
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

229.0.0.3.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
                State: <Active Int>
                Age: 2:09
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

229.0.0.4.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
```

```
                Next-hop reference count: 113
                State: <Active Int>
                Age: 2:12
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

229.0.0.5.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
                State: <Active Int>
                Age: 9
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

229.0.0.6.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
                State: <Active Int>
                Age: 2:15
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

229.0.0.7.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
                State: <Active Int>
                Age: 2:15
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

229.0.0.8.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
                State: <Active Int>
                Age: 2:15
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

229.0.0.9.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
                State: <Active Int>
                Age: 2:14
                Task: MC
                Announcement bits (1): 0-KRT
                AS path: I

229.0.0.10.11.11.11.100.3.13.0.0/80 (1 entry, 1 announced)
        *Multicast Preference: 180
                Next hop type: Multicast (IPv4)
                Next-hop reference count: 113
```

```
State: <Active Int>
Age: 2:13
Task: MC
Announcement bits (1): 0-KRT
AS path: I
```

show route source-gateway

Syntax
show route source-gateway *address*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show route source-gateway *address*
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display the entries in the routing table that were learned from a particular address. The **Source** field in the **show route detail** command output lists the source for each route, if known.

Options
brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

address—IP address of the system.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route source-gateway on page 2018
show route source-gateway detail on page 2019
show route source-gateway extensive on page 2021

Output Fields
For information about output fields, see the output field tables for the **show route** command, the **show route detail** command, the **show route extensive** command, or the **show route terse** command.

show route source-gateway

```
user@host> show route source-gateway 10.255.70.103
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

```
green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:1:3:1/96
                   *[BGP/170] 12:12:24, localpref 100, from 10.255.70.103
                      AS path: I
                    > via so-0/3/0.0, label-switched-path green-r1-r3

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:2:3:1/96
                   *[BGP/170] 12:12:24, localpref 0, from 10.255.70.103
                      AS path: I
                    > via so-0/3/0.0, label-switched-path green-r1-r3

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.255.70.103:1:3:1/96
                   *[BGP/170] 12:12:24, localpref 100, from 10.255.70.103
                      AS path: I
                    > via so-0/3/0.0, label-switched-path green-r1-r3

10.255.70.103:2:3:1/96
                   *[BGP/170] 12:12:24, localpref 0, from 10.255.70.103
                      AS path: I
                    > via so-0/3/0.0, label-switched-path green-r1-r3
```

show route source-gateway detail

```
user@host> show route source-gateway 10.255.70.103 detail
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete
green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

Restart Complete
10.255.70.103:1:3:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.70.103:1
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS:    69 Peer AS:    69
```

```
                    Age: 12:14:00   Metric2: 1
                    Task: BGP_69.10.255.70.103+179
                    Announcement bits (1): 0-green-l2vpn
                    AS path: I
                    Communities: target:11111:1 Layer2-info: encaps:VPLS,
                    control flags:, mtu: 0
                    Label-base: 800008, range: 8
                    Localpref: 100
                    Router ID: 10.255.70.103
                    Primary Routing Table bgp.l2vpn.0

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:2:3:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-1
                Route Distinguisher: 10.255.70.103:2
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 12:14:00   Metric2: 1
                Task: BGP_69.10.255.70.103+179
                Announcement bits (1): 0-red-l2vpn
                AS path: I
                Communities: target:11111:2 Layer2-info: encaps:VPLS,
                control flags:Site-Down, mtu: 0
                Label-base: 800016, range: 8
                Localpref: 0
                Router ID: 10.255.70.103
                Primary Routing Table bgp.l2vpn.0

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:1:3:1/96 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.70.103:1
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 12:14:00   Metric2: 1
                Task: BGP_69.10.255.70.103+179
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS, control
flags:, mtu: 0
                Label-base: 800008, range: 8
                Localpref: 100
                Router ID: 10.255.70.103
                Secondary Tables: green.l2vpn.0
10.255.70.103:2:3:1/96 (1 entry, 0 announced)
        *BGP    Preference: 170/-1
                Route Distinguisher: 10.255.70.103:2
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
```

```
                        Indirect next hop: 2 no-forward
                        State: <Active Int Ext>
                        Local AS:    69 Peer AS:    69
                        Age: 12:14:00   Metric2: 1
                        Task: BGP_69.10.255.70.103+179
                        AS path: I
                        Communities: target:11111:2 Layer2-info: encaps:VPLS,
                        control flags:Site-Down,
                        mtu: 0
                        Label-base: 800016, range: 8
                        Localpref: 0
                        Router ID: 10.255.70.103
                        Secondary Tables: red.l2vpn.0
```

show route source-gateway extensive

```
user@host> show route source-gateway 10.255.70.103 extensive
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
10.255.70.103:1:3:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.70.103:1
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 12:15:24   Metric2: 1
                Task: BGP_69.10.255.70.103+179
                Announcement bits (1): 0-green-l2vpn
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0
                Label-base: 800008, range: 8
                Localpref: 100
                Router ID: 10.255.70.103
                Primary Routing Table bgp.l2vpn.0

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:2:3:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-1
                Route Distinguisher: 10.255.70.103:2
                Next-hop reference count: 7
```

```
                    Source: 10.255.70.103
                    Protocol next hop: 10.255.70.103
                    Indirect next hop: 2 no-forward
                    State: <Secondary Active Int Ext>
                    Local AS:    69 Peer AS:    69
                    Age: 12:15:24   Metric2: 1
                    Task: BGP_69.10.255.70.103+179
                    Announcement bits (1): 0-red-l2vpn
                    AS path: I
                    Communities: target:11111:2 Layer2-info: encaps:VPLS,
                    control flags:Site-Down, mtu: 0
                    Label-base: 800016, range: 8
                    Localpref: 0
                    Router ID: 10.255.70.103
                    Primary Routing Table bgp.l2vpn.0

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

10.255.70.103:1:3:1/96 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.70.103:1
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 12:15:24   Metric2: 1
                Task: BGP_69.10.255.70.103+179
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0
                Label-base: 800008, range: 8
                Localpref: 100
                Router ID: 10.255.70.103
                Secondary Tables: green.l2vpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.255.70.103 Metric: 2
                        Indirect next hop: 2 no-forward
                        Indirect path forwarding next hops: 1
                Next hop:        via so-0/3/0.0 weight 0x1
                        10.255.70.103/32 Originating RIB: inet.3
                          Metric: 2                     Node path count: 1
                          Forwarding nexthops: 1
                                Nexthop: via so-0/3/0.0

10.255.70.103:2:3:1/96 (1 entry, 0 announced)
        *BGP    Preference: 170/-1
                Route Distinguisher: 10.255.70.103:2
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Active Int Ext>
                Local AS:    69 Peer AS:    69
                Age: 12:15:24   Metric2: 1
                Task: BGP_69.10.255.70.103+179
                AS path: I
                Communities: target:11111:2 Layer2-info: encaps:VPLS,
                control flags:Site-Down,
```

```
mtu: 0
Label-base: 800016, range: 8
Localpref: 0
Router ID: 10.255.70.103
Secondary Tables: red.l2vpn.0
Indirect next hops: 1
        Protocol next hop: 10.255.70.103 Metric: 2
        Indirect next hop: 2 no-forward
        Indirect path forwarding next hops: 1
Next hop:        via so-0/3/0.0 weight 0x1
        10.255.70.103/32 Originating RIB: inet.3
          Metric: 2                      Node path count: 1
          Forwarding nexthops: 1
                Nexthop: via so-0/3/0.0
```

show route summary

Syntax
show route summary
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show route summary

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display summary statistics about the entries in the routing table.

Options
none—Display summary statistics about the entries in the routing table.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route summary on page 2025

Output Fields
Table 252 on page 2024 lists the output fields for the **show route summary** command. Output fields are listed in the approximate order in which they appear.

Table 252: show route summary Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, **inet.0**).
destinations	Number of destinations for which there are routes in the routing table.
routes	Number of routes in the routing table: • **active**—Number of routes that are active. • **holddown**—Number of routes that are in the hold-down state before being declared inactive. • **hidden**—Number of routes not used because of routing policy.
Direct	Routes on the directly connected network.
Local	Local routes.
protocol-name	Name of the protocol from which the route was learned. For example, **OSPF**, **RSVP**, and **Static**.

show route summary

```
user@host> show route summary
Autonomous system number: 69
Router ID: 10.255.71.52
Maximum-ECMP: 32
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
              Direct:      6 routes,      5 active
               Local:      4 routes,      4 active
                OSPF:      5 routes,      4 active
              Static:      7 routes,      7 active
                IGMP:      1 routes,      1 active
                 PIM:      2 routes,      2 active

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
                RSVP:      2 routes,      2 active

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete
              Direct:      1 routes,      1 active

mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
                MPLS:      3 routes,      3 active
                VPLS:      4 routes,      2 active

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
Restart Complete
              Direct:      2 routes,      2 active
                 PIM:      2 routes,      2 active
                 MLD:      1 routes,      1 active

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
                 BGP:      2 routes,      2 active
               L2VPN:      2 routes,      2 active

red.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
Restart Complete
                 BGP:      2 routes,      2 active
               L2VPN:      1 routes,      1 active

bgp.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
                 BGP:      4 routes,      4 active
```

show route table

Syntax

show route table *routing-table-name*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)

show route table *routing-table-name*
<brief | detail | extensive | terse>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display the route entries in a particular routing table.

Options

brief | detail | extensive | terse—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

routing-table-name—Display information about a particular routing table (for example, **inet.0**) where policy-based export is currently enabled. (For information about the different types of routing tables, see the *Junos OS Routing Protocols Configuration Guide*.)

Required Privilege Level

view

Related Documentation

- **show route summary on page 2024**

List of Sample Output

show route table bgp.l2.vpn on page 2027
show route table bgp.l3vpn.0 on page 2027
show route table bgp.l3vpn.0 detail on page 2027
show route table inet.0 on page 2028
show route table inet6.0 on page 2029
show route table inet6.3 on page 2029
show route table l2circuit.0 on page 2029
show route table mpls on page 2030
show route table mpls extensive on page 2030
show route table mpls.0 on page 2030
show route table vpls_1 detail on page 2031
show route table vpn-a on page 2031
show route table vpn-a.mdt.0 on page 2031
show route table VPN-AB.inet.0 on page 2031
show route table VPN_blue.mvpn-inet6.0 on page 2032

Output Fields

For information about output fields, see the output field tables for the **show route** command, the **show route detail** command, the **show route extensive** command, or the **show route terse** command.

show route table bgp.l2.vpn

```
user@host> show route table bgp.l2.vpn
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.24.1:1:4:1/96
                   *[BGP/170] 01:08:58, localpref 100, from 192.168.24.1
                      AS path: I
                    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
```

show route table bgp.l3vpn.0

```
user@host> show route table bgp.l3vpn.0
      bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
      + = Active Route, - = Last Active, * = Both

      10.255.71.15:100:10.255.71.17/32
                         *[BGP/170] 00:03:59, MED 1, localpref 100, from
      10.255.71.15
                            AS path: I
                          > via so-2/1/0.0, Push 100020, Push 100011(top)
      10.255.71.15:200:10.255.71.18/32
                         *[BGP/170] 00:03:59, MED 1, localpref 100, from
      10.255.71.15
                            AS path: I
                          > via so-2/1/0.0, Push 100021, Push 100011(top)
```

show route table bgp.l3vpn.0 detail

```
user@host> show route table bgp.l3vpn.0 detail
bgp.l3vpn.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)

10.255.245.12:1:4.0.0.0/8 (1 entry, 1 announced)
  *BGP   Preference: 170/-101
         Route Distinguisher: 10.255.245.12:1
         Source: 10.255.245.12
         Next hop: 192.168.208.66 via fe-0/0/0.0, selected
         Label operation: Push 182449
         Protocol next hop: 10.255.245.12
         Push 182449
         Indirect next hop: 863a630 297
         State: <Active Int Ext>
         Local AS:    35 Peer AS:    35
         Age: 12:19      Metric2: 1
         Task: BGP_35.10.255.245.12+179
         Announcement bits (1): 0-BGP.0.0.0.0+179
         AS path: 30 10458 14203 2914 3356 I (Atomic) Aggregator: 3356 4.68.0.11

         Communities: 2914:420 target:11111:1 origin:56:78
         VPN Label: 182449
         Localpref: 100
          Router ID: 10.255.245.12

10.255.245.12:1:4.17.225.0/24 (1 entry, 1 announced)
  *BGP   Preference: 170/-101
         Route Distinguisher: 10.255.245.12:1
         Source: 10.255.245.12
         Next hop: 192.168.208.66 via fe-0/0/0.0, selected
         Label operation: Push 182465
         Protocol next hop: 10.255.245.12
         Push 182465
         Indirect next hop: 863a8f0 305
         State: <Active Int Ext>
         Local AS:    35 Peer AS:    35
         Age: 12:19      Metric2: 1
```

```
                Task: BGP_35.10.255.245.12+179
                Announcement bits (1): 0-BGP.0.0.0.0+179
AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496 6496 I
                       Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
                       VPN Label: 182465
                       Localpref: 100
                       Router ID: 10.255.245.12

10.255.245.12:1:4.17.226.0/23 (1 entry, 1 announced)
  *BGP   Preference: 170/-101
                Route Distinguisher: 10.255.245.12:1
                Source: 10.255.245.12
                Next hop: 192.168.208.66 via fe-0/0/0.0, selected
                Label operation: Push 182465
                Protocol next hop: 10.255.245.12
                Push 182465
                Indirect next hop: 86bd210 330
                State: <Active Int Ext>
                Local AS:    35 Peer AS:    35
                Age: 12:19      Metric2: 1
                Task: BGP_35.10.255.245.12+179
                Announcement bits (1): 0-BGP.0.0.0.0+179
                AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496

                6496 I
                Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
                VPN Label: 182465
                Localpref: 100
                Router ID: 10.255.245.12

10.255.245.12:1:4.17.251.0/24 (1 entry, 1 announced)
  *BGP   Preference: 170/-101
                Route Distinguisher: 10.255.245.12:1
                Source: 10.255.245.12
                Next hop: 192.168.208.66 via fe-0/0/0.0, selected
                Label operation: Push 182465
                Protocol next hop: 10.255.245.12
                Push 182465
                Indirect next hop: 86bd210 330
                State: <Active Int Ext>
                Local AS:    35 Peer AS:    35
                Age: 12:19      Metric2: 1
                Task: BGP_35.10.255.245.12+179
                Announcement bits (1): 0-BGP.0.0.0.0+179
                AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496

                6496 I
                Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
                VPN Label: 182465
                Localpref: 100
```

show route table inet.0

```
user@host> show route table inet.0
inet.0: 12 destinations, 12 routes (11 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 00:51:57
                    > to 111.222.5.254 via fxp0.0
1.0.0.1/32         *[Direct/0] 00:51:58
                    > via at-5/3/0.0
1.0.0.2/32         *[Local/0] 00:51:58
                      Local
```

```
12.12.12.21/32     *[Local/0] 00:51:57
                      Reject
13.13.13.13/32     *[Direct/0] 00:51:58
                    > via t3-5/2/1.0
13.13.13.14/32     *[Local/0] 00:51:58
                      Local
13.13.13.21/32     *[Local/0] 00:51:58
                      Local
13.13.13.22/32     *[Direct/0] 00:33:59
                    > via t3-5/2/0.0
127.0.0.1/32        [Direct/0] 00:51:58
                    > via lo0.0
111.222.5.0/24     *[Direct/0] 00:51:58
                    > via fxp0.0
111.222.5.81/32    *[Local/0] 00:51:58
                      Local
```

show route table inet6.0

```
user@host> show route table inet6.0
inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Route, * = Both

fec0:0:0:3::/64 *[Direct/0] 00:01:34
>via fe-0/1/0.0

fec0:0:0:3::/128 *[Local/0] 00:01:34
>Local

fec0:0:0:4::/64 *[Static/5] 00:01:34
>to fec0:0:0:3::ffff via fe-0/1/0.0
```

show route table inet6.3

```
user@router> show route table inet6.3
inet6.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::10.255.245.195/128
                   *[LDP/9] 00:00:22, metric 1
                    > via so-1/0/0.0
::10.255.245.196/128
                   *[LDP/9] 00:00:08, metric 1
                    > via so-1/0/0.0, Push 100008
```

show route table l2circuit.0

```
user@host> show route table l2circuit.0
l2circuit.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.195:NoCtrlWord:1:1:Local/96
                   *[L2CKT/7] 00:50:47
                    > via so-0/1/2.0, Push 100049
                      via so-0/1/3.0, Push 100049
10.1.1.195:NoCtrlWord:1:1:Remote/96
                   *[LDP/9] 00:50:14
                      Discard
10.1.1.195:CtrlWord:1:2:Local/96
                   *[L2CKT/7] 00:50:47
                    > via so-0/1/2.0, Push 100049
                      via so-0/1/3.0, Push 100049
10.1.1.195:CtrlWord:1:2:Remote/96
```

```
                        *[LDP/9] 00:50:14
                           Discard
```

show route table mpls

```
user@host> show route table mpls
mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                  *[MPLS/0] 00:13:55, metric 1
                      Receive
1                  *[MPLS/0] 00:13:55, metric 1
                      Receive
2                  *[MPLS/0] 00:13:55, metric 1
                      Receive
1024               *[VPN/0] 00:04:18
                      to table red.inet.0, Pop
```

show route table mpls extensive

```
user@host> show route table mpls extensive
100000 (1 entry, 1 announced)
TSI:
KRT in-kernel 100000 /36 -> {so-1/0/0.0}
        *LDP    Preference:   9
                Next hop:  via so-1/0/0.0, selected
                Pop
                State: <Active Int>
                Age: 29:50      Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
                Prefixes bound to route: 10.0.0.194/32
```

show route table mpls.0

```
user@host> show route table mpls.0
mpls.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                  *[MPLS/0] 00:45:09, metric 1
                      Receive
1                  *[MPLS/0] 00:45:09, metric 1
                      Receive
2                  *[MPLS/0] 00:45:09, metric 1
                      Receive
100000             *[L2VPN/7] 00:43:04
                    > via so-0/1/0.1, Pop
100001             *[L2VPN/7] 00:43:03
                    > via so-0/1/0.2, Pop       Offset: 4
100002             *[LDP/9] 00:43:22, metric 1
                      via so-0/1/2.0, Pop
                    > via so-0/1/3.0, Pop
100002(S=0)        *[LDP/9] 00:43:22, metric 1
                      via so-0/1/2.0, Pop
                    > via so-0/1/3.0, Pop
100003             *[LDP/9] 00:43:22, metric 1
                    > via so-0/1/2.0, Swap 100002
                      via so-0/1/3.0, Swap 100002
100004             *[LDP/9] 00:43:16, metric 1
                      via so-0/1/2.0, Swap 100049
                    > via so-0/1/3.0, Swap 100049
so-0/1/0.1         *[L2VPN/7] 00:43:04
                    > via so-0/1/2.0, Push 100001, Push 100049(top)
                      via so-0/1/3.0, Push 100001, Push 100049(top)
so-0/1/0.2         *[L2VPN/7] 00:43:03
```

```
                                  via so-0/1/2.0, Push 100000, Push 100049(top) Offset: -4
                                > via so-0/1/3.0, Push 100000, Push 100049(top) Offset: -4
```

show route table vpls_1 detail

```
user@host> show route table vpls_1 detail
vpls_1.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

1.1.1.11:1000:1:1/96 (1 entry, 1 announced)
*L2VPN Preference: 170/-1
Receive table: vpls_1.l2vpn.0
Next-hop reference count: 2
State: <Active Int Ext>
Age: 4:29:47 Metric2: 1
Task: vpls_1-l2vpn
Announcement bits (1): 1-BGP.0.0.0.0+179
AS path: I
Communities: Layer2-info: encaps:VPLS, control flags:Site-Down
Label-base: 800000, range: 8, status-vector: 0xFF
```

show route table vpn-a

```
user@host> show route table vpn-a
vpn-a.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both
192.168.16.1:1:1:1/96
                   *[VPN/7] 05:48:27
                     Discard
192.168.24.1:1:2:1/96
                   *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
                      AS path: I
                    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
192.168.24.1:1:3:1/96
                   *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
                      AS path: I
                    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
```

show route table vpn-a.mdt.0

```
user@host> show route table vpn-a.mdt.0
vpn-a.mdt.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:1:0:10.255.14.216:232.1.1.1/144
                   *[MVPN/70] 01:23:05, metric2 1
                      Indirect
1:1:1:10.255.14.218:232.1.1.1/144
                   *[BGP/170] 00:57:49, localpref 100, from 10.255.14.218
                      AS path: I
                    > via so-0/0/0.0, label-switched-path r0e-to-r1
1:1:2:10.255.14.217:232.1.1.1/144
                   *[BGP/170] 00:57:49, localpref 100, from 10.255.14.217
                      AS path: I
> via so-0/0/1.0, label-switched-path r0-to-r2
```

show route table VPN-AB.inet.0

```
user@host> show route table VPN-AB.inet.0
      VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
      + = Active Route, - = Last Active, * = Both

      10.39.1.0/30       *[OSPF/10] 00:07:24, metric 1
                          > via so-7/3/1.0
      10.39.1.4/30       *[Direct/0] 00:08:42
                          > via so-5/1/0.0
      10.39.1.6/32       *[Local/0] 00:08:46
```

```
                        Local
10.255.71.16/32    *[Static/5] 00:07:24
                    > via so-2/0/0.0
10.255.71.17/32    *[BGP/170] 00:07:24, MED 1, localpref 100, from
10.255.71.15
                      AS path: I
                     > via so-2/1/0.0, Push 100020, Push 100011(top)
10.255.71.18/32    *[BGP/170] 00:07:24, MED 1, localpref 100, from
10.255.71.15
                      AS path: I
                     > via so-2/1/0.0, Push 100021, Push 100011(top)
10.255.245.245/32  *[BGP/170] 00:08:35, localpref 100
                      AS path: 2 I
                     > to 10.39.1.5 via so-5/1/0.0
10.255.245.246/32  *[OSPF/10] 00:07:24, metric 1
                     > via so-7/3/1.0
```

show route table VPN_blue.mvpn-inet6.0

```
user@host> show route table VPN_blue.mvpn-inet6.0
vpn_blue.mvpn-inet6.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:10.255.2.202:65535:10.255.2.202/432
                   *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
                      AS path: I
                     > via so-0/1/3.0
1:10.255.2.203:65535:10.255.2.203/432
                   *[BGP/170] 00:02:37, localpref 100, from 10.255.2.203
                      AS path: I
                     > via so-0/1/0.0
1:10.255.2.204:65535:10.255.2.204/432
                   *[MVPN/70] 00:57:23, metric2 1
                      Indirect
5:10.255.2.202:65535:128:::192.168.90.2:128:ffff::1/432
                   *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
                      AS path: I
                     > via so-0/1/3.0
6:10.255.2.203:65535:65000:128:::10.12.53.12:128:ffff::1/432
                   *[PIM/105] 00:02:37
                      Multicast (IPv6)
7:10.255.2.202:65535:65000:128:::192.168.90.2:128:ffff::1/432
                   *[MVPN/70] 00:02:37, metric2 1
                      Indirect
```

show route terse

Syntax show route terse
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch) show route terse

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display a high-level summary of the routes in the routing table.

Options none—Display a high-level summary of the routes in the routing table.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level view

List of Sample Output **show route terse on page 2035**

Output Fields Table 253 on page 2033 describes the output fields for the **show route terse** command. Output fields are listed in the approximate order in which they appear.

Table 253: show route terse Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, **inet.0**).
number **destinations**	Number of destinations for which there are routes in the routing table.
number **routes**	Number of routes in the routing table and total number of routes in the following states: • **active** (routes that are active) • **holddown** (routes that are in the pending state before being declared inactive) • **hidden** (routes that are not used because of a routing policy)
route key	Key for the state of the route: • **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • **-** —A hyphen indicates the last active route. • *****—An asterisk indicates that the route is both the active and the last active route. An asterisk before a **to** line indicates the best subpath to the route.
A	Active route. An asterisk (*) indicates this is the active route.
Destination	Destination of the route.

Table 253: show route terse Output Fields *(continued)*

Field Name	Field Description
P	Protocol through which the route was learned: • **A**—Aggregate • **B**—BGP • **C**—CCC • **D**—Direct • **G**—GMPLS • **I**—IS-IS • **L**—L2CKT, L2VPN, LDP, Local • **K**—Kernel • **M**—MPLS, MSDP • **O**—OSPF • **P**—PIM • **R**—RIP, RIPng • **S**—Static • **T**—Tunnel
Prf	Preference value of the route. In every routing metric except for the BGP **LocalPref** attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the **LocalPref** value in the **Preference2** field. For example, if the **LocalPref** value for Route 1 is 100, the **Preference2** value is -101. If the **LocalPref** value for Route 2 is 155, the **Preference2** value is -156. Route 2 is preferred because if has a higher **LocalPref** value and a lower **Preference2** value.
Metric 1	First metric value in the route. For routes learned from BGP, this is the MED metric.
Metric 2	Second metric value in the route. For routes learned from BGP, this is the IGP metric.
Next hop	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **?**—Incomplete; typically, the AS path was aggregated.

show route terse

```
user@host> show route terse
inet.0: 12 destinations, 12 routes (11 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf Metric 1   Metric 2    Next hop        AS path
* 0.0.0.0/0          S   5                       >111.222.5.254
* 1.0.0.1/32         D   0                       >at-5/3/0.0
* 1.0.0.2/32         L   0                        Local
* 12.12.12.21/32     L   0                        Reject
* 13.13.13.13/32     D   0                       >t3-5/2/1.0
* 13.13.13.14/32     L   0                        Local
* 13.13.13.21/32     L   0                        Local
* 13.13.13.22/32     D   0                       >t3-5/2/0.0
  127.0.0.1/32       D   0                       >lo0.0
* 111.222.5.0/24     D   0                       >fxp0.0
* 111.222.5.81/32    L   0                        Local
* 224.0.0.5/32       O  10          1             MultiRecv
```

show vrrp

Syntax
show vrrp
<brief | detail | extensive | summary>
<interface *interface-name*>
<track interfaces>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display information and status about VRRP groups.

Options
none—(Same as brief) Display brief status information about all VRRP interfaces.

brief | detail | extensive | summary—(Optional) Display the specified level of output.

interface *interface-name* —(Optional) Display information and status about the specified VRRP interface.

track interfaces—(Optional) Display information and status about VRRP track interfaces.

Required Privilege Level
view

Related Documentation
- Configuring VRRP for IPv6 (CLI Procedure) on page 1452

List of Sample Output
show vrrp on page 2041
show vrrp brief on page 2041
show vrrp detail (IPv6) on page 2041
show vrrp detail (Route Track) on page 2041
show vrrp extensive on page 2041
show vrrp interface on page 2043
show vrrp summary on page 2044
show vrrp track detail on page 2044
show vrrp track summary on page 2044

Output Fields
Table 254 on page 2036 lists the output fields for the **show vrrp** command. Output fields are listed in the approximate order in which they appear.

Table 254: show vrrp Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the logical interface.	**none, brief, extensive, summary**
Interface index	Physical interface index number, which reflects its initialization sequence.	**extensive**
Groups	Total number of VRRP groups configured on the interface.	**extensive**
Active	Total number of VRRP groups that are active (that is, whose interface state is either up or down).	**extensive**

Table 254: show vrrp Output Fields *(continued)*

Field Name	Field Description	Level of Output
Interface VRRP PDU statistics	Nonerrored statistics for the logical interface: • **Advertisement sent**—Number of VRRP advertisement protocol data units (PDUs) that the interface has transmitted. • **Advertisement received**—Number of VRRP advertisement PDUs received by the interface. • **Packets received**—Number of VRRP packets received for VRRP groups on the interface. • **No group match received**—Number of VRRP packets received for VRRP groups that do not exist on the interface.	**extensive**
Interface VRRP PDU error statistics	Errored statistics for the logical interface: • **Invalid IPAH next type received**—Number of packets received that use the IP Authentication Header protocol (IPAH) and that do not encapsulate VRRP packets. • **Invalid VRRP ttl value received**—Number of packets received whose IP time-to-live (TTL) value is not 255. • **Invalid VRRP version received**—Number of packets received whose VRRP version is not 2. • **Invalid VRRP pdu type received**—Number of packets received whose VRRP PDU type is not 1. • **Invalid VRRP authentication type received**—Number of packets received whose VRRP authentication is not none, simple, or md5. • **Invalid VRRP IP count received**—Number of packets received whose VRRP IP count exceeds 8. • **Invalid VRRP checksum received**—Number of packets received whose VRRP checksum does not match the calculated value.	**extensive**
Physical interface	Name of the physical interface.	**detail, extensive**
Unit	Logical unit number.	All levels
Address	Address of the physical interface.	**none, brief, detail, extensive**
Index	Physical interface index number, which reflects its initialization sequence.	**detail, extensive**
SNMP ifIndex	SNMP index number for the physical interface.	**detail, extensive**
VRRP-Traps	Status of VRRP traps: **Enabled** or **Disabled**.	**detail, extensive**
Type and Address	Identifier for the address and the address itself: • **lcl**—Configured local interface address. • **mas**—Address of the master virtual router. This address is displayed only when the local interface is acting as a backup router. • **vip**—Configured virtual IP addresses.	**none, brief, summary**

Table 254: show vrrp Output Fields *(continued)*

Field Name	Field Description	Level of Output
Interface state or **Int state**	State of the physical interface: • **down**—The device is present and the link is unavailable. • **not present**—The interface is configured, but no physical device is present. • **unknown**—The VRRP process has not had time to query the kernel about the state of the interface. • **up**—The device is present and the link is established.	**none, brief, extensive, summary**
Group	VRRP group number.	**none, brief, extensive, summary**
State	VRRP state: • **backup**—The interface is acting as the backup router interface. • **bringup**—VRRP is just starting, and the physical device is not yet present. • **idle**—VRRP is configured on the interface and is disabled. This can occur when VRRP is first enabled on an interface whose link is established. • **initializing**—VRRP is initializing. • **master**—The interface is acting as the master router interface. • **transition**—The interface is changing between being the backup and being the master router.	**extensive**
Priority	Configured VRRP priority for the interface.	**detail, extensive**
Advertisement interval	Configured VRRP advertisement interval.	**detail, extensive**
Authentication type	Configured VRRP authentication type: **none**, **simple**, or **md5**.	**detail, extensive**
Preempt	Whether preemption is allowed on the interface: **yes** or **no**.	**detail, extensive**
Accept-data mode	Whether the interface is configured to accept packets destined for the virtual IP address: **yes** or **no**.	**detail, extensive**
VIP count	Number of virtual IP addresses that have been configured on the interface.	**detail, extensive**
VIP	List of virtual IP addresses configured on the interface.	**detail, extensive**
Advertisement timer	Time until the advertisement timer expires.	**detail, extensive**
Master router	IP address of the interface that is acting as the master. If the VRRP interface is down, the output is **N/A**.	**detail, extensive**
Virtual router uptime	Time that the virtual router has been up.	**detail, extensive**
Master router uptime	Time that the master router has been up.	**detail, extensive**

Table 254: show vrrp Output Fields *(continued)*

Field Name	Field Description	Level of Output
Virtual MAC	MAC address associated with the virtual IP address.	**detail, extensive**
Tracking	Whether tracking is **enabled** or **disabled**.	**detail, extensive**
Current priority	Current operational priority for being the VRRP master.	**detail, extensive**
Configured priority	Configured base priority for being the VRRP master.	**detail, extensive**
Priority hold-time	Minimum time interval, in seconds, between successive changes to the current priority. **Disabled** indicates no minimum interval.	**detail, extensive**
Remaining-time	(**track** option only) Displays the time remaining in the priority hold-time interval.	**detail**
Interface tracking	Whether interface tracking is enabled or disabled. When enabled, the output also displays the number of tracked interfaces.	**detail extensive**
Interface/Tracked interface	Name of the tracked interface.	**detail extensive**
Int state/Interface state	Current operational state of the tracked interface: **up** or **down**.	**detail, extensive**
Int speed/Speed	Current operational speed, in bits per second, of the tracked interface.	**detail, extensive**
Incurred priority cost	Operational priority cost incurred due to the state and speed of this tracked interface. This cost is applied to the configured priority to obtain the current priority.	**detail, extensive**
Threshold	Speed below which the corresponding priority cost is incurred. In other words, when the speed of the interface drops below the threshold speed, the corresponding priority cost is incurred. An entry of **down** means that the corresponding priority cost is incurred when the interface is down.	**detail, extensive**
Route tracking	Whether route tracking is enabled or disabled. When enabled, the output also displays the number of tracked routes.	**detail, extensive**
Route count	The number of routes being tracked.	**detail, extensive**
Route	The IP address of the route being tracked.	**detail, extensive**
VRF name	The VPN routing and forwarding (VRF) routing instance that the tracked route is in.	**detail, extensive**
Route state	The state of the route being tracked: **up**, **down**, or **unknown**.	**detail, extensive**
Priority cost	Configured priority cost. This value is incurred when the interface speed drops below the corresponding threshold or when the tracked route goes down.	**detail, extensive**

Table 254: show vrrp Output Fields *(continued)*

Field Name	Field Description	Level of Output
Active	Whether the threshold is active (*). If the threshold is active, the corresponding priority cost is incurred.	**detail, extensive**
Group VRRP PDU statistics	Number of VRRP advertisements sent and received by the group.	**extensive**
Group VRRP PDU error statistics	Errored statistics for the VRRP group: • **Bad authentication type received**—Number of VRRP PDUs received with an invalid authentication type. The received authentication can be **none**, **simple**, or **md5** and must be the same for all routers in the VRRP group. • **Bad password received**—Number of VRRP PDUs received with an invalid key (password). The password for simple authentication must be the same for all routers in the VRRP group • **Bad MD5 digest received**—Number of VRRP PDUs received for which the MD5 digest computed from the VRRP PDU differs from the digest expected by the VRRP instance configured on the router. • **Bad advertisement timer received**—Number of VRRP PDUs received with an advertisement time interval that is inconsistent with the one in use among the routers in the VRRP group. • **Bad VIP count received**—Number of VRRP PDUs whose virtual IP address counts differ from the count that has been configured on the VRRP instance. • **Bad VIPADDR received**—Number of VRRP PDUs whose virtual IP addresses differ from the list of virtual IP addresses configured on the VRRP instance.	**extensive**
Group state transition statistics	State transition statistics for the VRRP group: • **Idle to master transitions**—Number of times that the VRRP instance transitioned from the idle state to the master state. • **Idle to backup transitions**—Number of times that the VRRP instance transitioned from the idle state to the backup state. • **Backup to master transitions**—Number of times that the VRRP instance transitioned from the backup state to the master state. • **Master to backup transitions**—Number of times that the VRRP instance transitioned from the master state to the backup state.	**extensive**
VR state	VRRP information: • **backup**—The interface is acting as the backup router interface. • **bringup**—VRRP is just starting, and the physical device is not yet present. • **idle**—VRRP is configured on the interface and is disabled. This can occur when VRRP is first enabled on an interface whose link is established. • **initializing**—VRRP is initializing. • **master**—The interface is acting as the master router interface. • **transition**—The interface is changing between being the backup and being the master router.	**none, brief**
Timer	VRRP timer information: • **A**—Time, in seconds, until the advertisement timer expires. • **D**—Time, in seconds, until the Master is Dead timer expires.	**none, brief**

show vrrp

```
user@host> show vrrp
Interface         State        Group    VR state    Timer   Type  Address
ge-0/0/0.121      up               1    master      A 1.052  lcl   gec0::12:1:1:1

                                                              vip   ge80::12:1:1:99

                                                              vip   gec0::12:1:1:99
ge-0/0/2.131      up               1    master      A 0.364  lcl   gec0::13:1:1:1

                                                              vip   ge80::13:1:1:99

                                                              vip   gec0::13:1:1:99
```

show vrrp brief

The output for the **show vrrp brief** command is identical to that for the **show vrrp** command. For sample output, see **show vrrp on page 2041**.

show vrrp detail (IPv6)

```
user@host> show vrrp detail
Physical interface: ge-0/0/0, Unit: 121, Vlan-id: 212, Address: gec0::12:1:1:1/120

  Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled
  Interface state: up, Group: 1, State: master
  Priority: 200, Advertisement interval: 1, Authentication type: none
  Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::12:1:1:99,
gec0::12:1:1:99
  Advertisement timer: 1.121s, Master router: ge80::12:1:1:1
  Virtual router uptime: 00:03:47, Master router uptime: 00:03:41
  Virtual MAC: 00:00:5e:00:02:01
  Tracking: disabled

Physical interface: ge-0/0/2, Unit: 131, Vlan-id: 213, Address: gec0::13:1:1:1/120

  Index: 69, SNMP ifIndex: 47, VRRP-Traps: enabled
  Interface state: up, Group: 1, State: master
  Priority: 200, Advertisement interval: 1, Authentication type: none
  Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::13:1:1:99,
gec0::13:1:1:99
  Advertisement timer: 0.327s, Master router: ge80::13:1:1:1
  Virtual router uptime: 00:03:47, Master router uptime: 00:03:41
  Virtual MAC: 00:00:5e:00:02:01
  Tracking: disabled
```

show vrrp detail (Route Track)

```
user@host> show vrrp detail
Physical interface: ge-1/1/0, Unit: 0, Address: 30.30.30.30/24
  Index: 67, SNMP ifIndex: 379, VRRP-Traps: enabled
  Interface state: up, Group: 100, State: master
  Priority: 150, Advertisement interval: 1, Authentication type: none
  Preempt: yes, Accept-data mode: no, VIP count: 1, VIP: 30.30.30.100
  Advertisement timer: 1.218s, Master router: 30.30.30.30
  Virtual router uptime: 00:04:28, Master router uptime: 00:00:13
  Virtual MAC: 00:00:5e:00:01:64
  Tracking: enabled
    Current priority: 150, Configured priority: 150
    Priority hold-time: disabled
    Interface tracking: disabled
    Route tracking: enabled, Route count: 1
      Route               VRF name       Route state   Priority cost
      192.168.40.0/22     default        up            30
```

show vrrp extensive

```
user@host> show vrrp extensive
```

```
Interface: ge-0/0/0.121, Interface index: 67, Groups: 1, Active : 1
  Interface VRRP PDU statistics
    Advertisement sent                      :          188
    Advertisement received                  :            0
    Packets received                        :            0
    No group match received                 :            0
  Interface VRRP PDU error statistics
    Invalid IPAH next type received         :            0
    Invalid VRRP TTL value received         :            0
    Invalid VRRP version received           :            0
    Invalid VRRP PDU type received          :            0
    Invalid VRRP authentication type received:           0
    Invalid VRRP IP count received          :            0
    Invalid VRRP checksum received          :            0

Physical interface: ge-0/0/0, Unit: 121, Vlan-id: 212, Address: gec0::12:1:1:1/120

  Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled
  Interface state: up, Group: 1, State: master
  Priority: 200, Advertisement interval: 1, Authentication type: none
  Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::12:1:1:99,
gec0::12:1:1:99
  Advertisement timer: 1.034s, Master router: ge80::12:1:1:1
  Virtual router uptime: 00:04:04, Master router uptime: 00:03:58
  Virtual MAC: 00:00:5e:00:02:01
  Tracking: disabled
  Group VRRP PDU statistics
    Advertisement sent              :          188
    Advertisement received          :            0
  Group VRRP PDU error statistics
    Bad authentication type received:            0
    Bad password received           :            0
    Bad MD5 digest received         :            0
    Bad advertisement timer received:            0
    Bad VIP count received          :            0
    Bad VIPADDR received            :            0
  Group state transition statistics
    Idle to master transitions      :            0
    Idle to backup transitions      :            1
    Backup to master transitions    :            1
    Master to backup transitions    :            0

Interface: ge-0/0/2.131, Interface index: 69, Groups: 1, Active : 1
  Interface VRRP PDU statistics
    Advertisement sent                      :          186
    Advertisement received                  :            0
    Packets received                        :            0
    No group match received                 :            0
  Interface VRRP PDU error statistics
    Invalid IPAH next type received         :            0
    Invalid VRRP TTL value received         :            0
    Invalid VRRP version received           :            0
    Invalid VRRP PDU type received          :            0
    Invalid VRRP authentication type received:           0
    Invalid VRRP IP count received          :            0
    Invalid VRRP checksum received          :            0

Physical interface: ge-0/0/2, Unit: 131, Vlan-id: 213, Address: gec0::13:1:1:1/120

  Index: 69, SNMP ifIndex: 47, VRRP-Traps: enabled
  Interface state: up, Group: 1, State: master
```

```
  Priority: 200, Advertisement interval: 1, Authentication type: none
  Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::13:1:1:99,
gec0::13:1:1:99
  Advertisement timer: 0.396s, Master router: ge80::13:1:1:1
  Virtual router uptime: 00:04:04, Master router uptime: 00:03:58
  Virtual MAC: 00:00:5e:00:02:01
  Tracking: disabled
  Group VRRP PDU statistics
    Advertisement sent              :          186
    Advertisement received          :            0
  Group VRRP PDU error statistics
    Bad authentication type received:            0
    Bad password received           :            0
    Bad MD5 digest received         :            0
    Bad advertisement timer received:            0
    Bad VIP count received          :            0
    Bad VIPADDR received            :            0
  Group state transition statistics
    Idle to master transitions      :            0
    Idle to backup transitions      :            1
    Backup to master transitions    :            1
  Master to backup transitions    :           0
```

show vrrp interface

```
user@host> show vrrp interface
Interface: ge-0/0/0.121, Interface index: 67, Groups: 1, Active : 1
  Interface VRRP PDU statistics
    Advertisement sent                      :          205
    Advertisement received                  :            0
    Packets received                        :            0
    No group match received                 :            0
  Interface VRRP PDU error statistics
    Invalid IPAH next type received         :            0
    Invalid VRRP TTL value received         :            0
    Invalid VRRP version received           :            0
    Invalid VRRP PDU type received          :            0
    Invalid VRRP authentication type received:           0
    Invalid VRRP IP count received          :            0
    Invalid VRRP checksum received          :            0

Physical interface: ge-0/0/0, Unit: 121, Vlan-id: 212, Address: gec0::12:1:1:1/120

  Index: 67, SNMP ifIndex: 45, VRRP-Traps: enabled
  Interface state: up, Group: 1, State: master
  Priority: 200, Advertisement interval: 1, Authentication type: none
  Preempt: yes, Accept-data mode: no, VIP count: 2, VIP: ge80::12:1:1:99,
gec0::12:1:1:99
  Advertisement timer: 0.789s, Master router: ge80::12:1:1:1
  Virtual router uptime: 00:04:26, Master router uptime: 00:04:20
  Virtual MAC: 00:00:5e:00:02:01
  Tracking: disabled
  Group VRRP PDU statistics
    Advertisement sent              :          205
    Advertisement received          :            0
  Group VRRP PDU error statistics
    Bad authentication type received:            0
    Bad password received           :            0
    Bad MD5 digest received         :            0
    Bad advertisement timer received:            0
    Bad VIP count received          :            0
    Bad VIPADDR received            :            0
  Group state transition statistics
```

```
    Idle to master transitions      :           0
    Idle to backup transitions      :           1
    Backup to master transitions    :           1
    Master to backup transitions    :           0
```

show vrrp summary

```
user@host> show vrrp summary
Interface     State       Group   VR state      Type   Address
ge-4/1/0.0    up              1   backup        lcl    10.57.0.2
                                                vip    10.57.0.100
```

show vrrp track detail

```
user@host> show vrrp track detail
Tracked interface: ae1.211
  State: up, Speed: 400m
  Incurred priority cost: 0
  Threshold   Priority cost   Active
  400m        10
  300m        60
  200m        110
  100m        160
  down        190
  Tracking VRRP interface: ae0.210, Group: 1
    VR State: master
    Current priority: 200, Configured priority: 200
    Priority hold-time: disabled,    Remaining-time: 50.351
```

show vrrp track summary

```
user@host> show vrrp track summary
Track if      State     Speed       VRRP if   Group   VR State   Current priority
ae1.211       up        400m        ae0.210    1      master     200
```

PART 16

IGMP Snooping and Multicast

CHAPTER 75

Understanding IGMP Snooping and Multicast

IGMP Snooping on J-EX Series Switches Overview

Internet Group Management Protocol (IGMP) snooping regulates multicast traffic in a switched network. With IGMP snooping enabled, a LAN switch monitors the IGMP transmissions between a host (a network device) and a multicast router, keeping track of the multicast groups and associated member interfaces. The switch uses that information to make intelligent multicast-forwarding decisions and forward traffic to the intended destination interfaces. J-EX Series Switches support IGMPv1, IGMPv2, and IGMPv3.

For details on IGMPv1, IGMPv2, and IGMPv3, see the following standards:

- For IGMPv1, see RFC 1112, *Host extensions for IP multicasting* at http://www.faqs.org/rfcs/rfc1112.html.
- For IGMPv2, see RFC 2236, *Internet Group Management Protocol, Version 2* at http://www.faqs.org/rfcs/rfc2236.html.
- For IGMPv3, see RFC 3376, *Internet Group Management Protocol, Version 3* at http://www.faqs.org/rfcs/rfc3376.html.

This IGMP snooping topic covers:

How IGMP Snooping Works

A J-EX Series switch usually learns *unicast* media access control (MAC) addresses by checking the source address field of the frames it receives. However, a *multicast* MAC

address can never be the source address for a packet. As a result, the switch floods multicast traffic on the VLAN, consuming significant amounts of bandwidth.

IGMP snooping regulates multicast traffic on a VLAN to avoid flooding. When IGMP snooping is enabled, the switch intercepts IGMP packets and uses the content of the packets to build a multicast cache table. The cache table is a database of multicast groups and their corresponding member ports. The cache table is then used to regulate multicast traffic on the VLAN.

When the switch receives multicast packets, it uses the cache table to selectively forward the packets only to the ports that are members of the destination multicast group. Figure 41 on page 2048 shows an example of IGMP traffic flow with IGMP snooping enabled.

Figure 41: IGMP Traffic Flow with IGMP Snooping Enabled

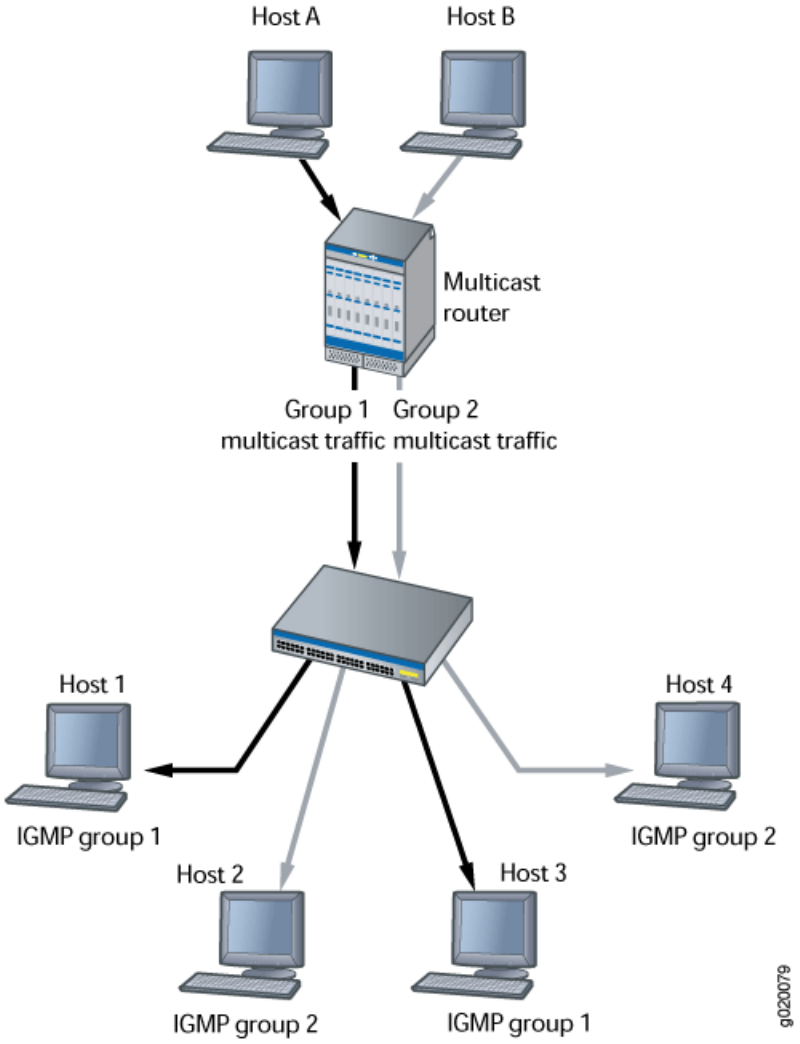

How IGMP Snooping Works with Routed VLAN Interfaces

Switches send traffic to hosts that are part of the same broadcast domain, but routers are needed to route traffic from one broadcast domain to another. Switches use a routed

VLAN interface (RVI) to perform these routing functions. IGMP snooping works with Layer 2 interfaces and RVIs to regulate multicast traffic in a switched network.

When a switch receives a multicast packet, the Packet Forwarding Engines in the switch perform an IP multicast lookup on the multicast packet to determine how to forward the packet to its local ports. From the results of the IP multicast lookup, each Packet Forwarding Engine extracts a list of Layer 3 interfaces (which can include VLAN interfaces) that have ports local to the Packet Forwarding Engine. If an RVI is part of this list, the switch provides a bridge multicast group ID for each RVI to the Packet Forwarding Engine.

A bridge multicast ID is assigned to direct Layer 3 interfaces and to RVIs. For VLANs that include multicast receivers, the bridge multicast ID includes a sub-next-hop ID. The sub-next-hop ID identifies the multicast Layer 2 interfaces in that VLAN that are interested in receiving the multicast stream. The switch ultimately assigns a next hop after it does a route lookup. The next hop includes all direct Layer 3 interfaces and RVIs. The Packet Forwarding Engine then forwards multicast traffic to the bridge multicast ID that includes all Layer 3 interfaces and RVIs that are multicast receivers for a given multicast group.

Figure 42 on page 2050 shows how multicast traffic is forwarded on a multilayer switch. In this illustration, multicast traffic is coming in through the **xe-0/1/0.0** interface. A multicast group has been formed by the Layer 3 interface **ge-0/0/2.0**, **vlan.0**, and **vlan.1**. The **ge-2/0/0.0** interface is a common trunk interface that belongs to both **vlan.0** and **vlan.1**. The letter “R” next to an interface name in the illustration indicates that a multicast receiver host is associated with that interface.

NOTE: Traffic sent to an access interface is untagged; traffic sent to a trunk interface is tagged. For more information on VLAN tagging, see “Understanding Bridging and VLANs on J-EX Series Switches” on page 1041.

Figure 42: IGMP Traffic Flow with Routed VLAN Interfaces

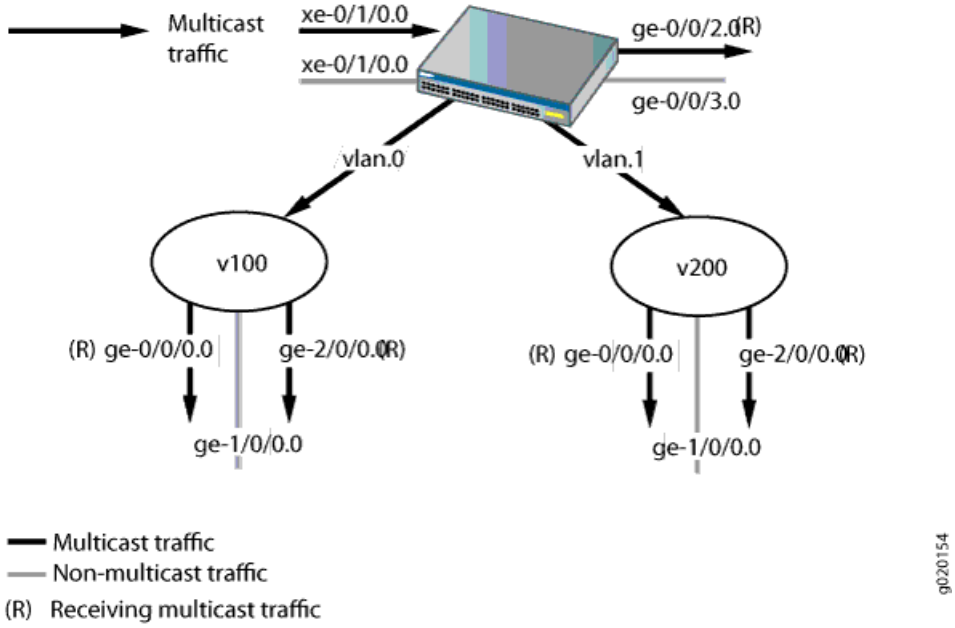

Table 255 on page 2050 shows the bridge multicast IDs and next hops that are created. The term **subnh** refers to a sub-next hop. The Packet Forwarding Engine will forward multicast traffic to bridge multicast ID9.

Table 255: Bridge Multicast IDs and Next Hops

ID Number	Type of Next Hop	Next Hop	Tag Information
ID1	RHN_UNICAST	ge-0/0/0.0	tag=off
ID2	RHN_UNICAST	ge-2/0/0.0	tag=on
ID3	RHN_FLOOD	[ID1, ID2]	
ID4	RHN_UNICAST	ge-0/0/1.0	tag=off
ID5	RHN_FLOOD	[ID4, ID2]	
ID6	RHN_UNICAST	vlan.0	subnh=ID3
ID7	RHN_UNICAST	VLAN.1	subnh=ID5
ID8	RHN_UNICAST	ge-0/0/2.0	
ID9	RHN_FLOOD	[ID6, ID7, ID8]	

How Hosts Join and Leave Multicast Groups

Hosts can join multicast groups in either of two ways:

- By sending an unsolicited IGMP join message to a multicast router that specifies the IP multicast group that the host is attempting to join.
- By sending an IGMP join message in response to a general query from a multicast router.

A multicast router continues to forward multicast traffic to a VLAN provided that at least one host on that VLAN responds to the periodic general IGMP queries. For a host to remain a member of a multicast group, therefore, it must continue to respond to the periodic general IGMP queries.

To leave a multicast group, a host can either not respond to the periodic general IGMP queries, which results in a "silent leave" (the only leave option for hosts connected to switches running IGMPv1), or send a group-specific IGMPv2 leave message.

NOTE: A host does not leave a group if its link goes down—for example, if a user disconnects from the port. The host remains a member of the group until group membership times out and a silent leave occurs. This means that if another user connects to the port before the silent leave occurs, the host resumes receiving the group multicast traffic until the silent leave, even though it never sent an IGMP join message.

IGMP Snooping Support for IGMPv3

IGMPv3 allows IGMP snooping to filter multicast streams based on the source address of the multicast stream. Junos OS for J-EX Series switches supports IGMPv3 packets that are in INCLUDE or EXCLUDE mode.

When a host sends an IGMPv3 INCLUDE report through a switch interface to indicate that it wants to receive a multicast stream from a source address, the switch adds the source address to the source list. In INCLUDE mode, the switch requests that packets be sent to the specified multicast address only from those IP source addresses listed in the source-list parameter. However, because J-EX Series switches do not support forwarding on a per-source basis, the switch merges all IGMPv3 reports for a VLAN to create a **(*,G,V)** route with the appropriate next hop. This next hop contains all the interfaces on the VLAN that are interested in group **G**.

When IGMP snooping for IGMPv3 is used with an RVI, the same **(*,G,V)** route is added to the snooping information in the RVI's output interface list (**olist**).

When a host sends an IGMPv3 EXCLUDE report, the host indicates that it wants to join a multicast group and receive packets for that group *except* from those IP source addresses in the source-list parameter. However, because J-EX Series switches do not support forwarding on a per-source basis, the switch ignores the source information and creates a **(*,G,V)** route. A host can also send an EXCLUDE report in which the source-list parameter is empty, which is known as an EXCLUDE NULL report. An EXCLUDE NULL

report indicates that the host wants to join the multicast group and receive packets from all sources. The switch creates a **(*, G,V)** route in this case also.

Related Documentation

- Understanding Multicast VLAN Registration on J-EX Series Switches on page 2052
- Example: Configuring IGMP Snooping on J-EX Series Switches on page 2055
- Configuring IGMP Snooping (CLI Procedure) on page 2063
- RFC 3171, *IANA Guidelines for IPv4 Multicast Address Assignments* at http://tools.ietf.org/html/rfc3171

Understanding Multicast VLAN Registration on J-EX Series Switches

Multicast VLAN registration (MVR) allows you to efficiently distribute IPTV multicast streams across an Ethernet ring-based Layer 2 network and reduce the amount of bandwidth consumed by this multicast traffic.

In a standard Layer 2 network, a multicast stream received on one VLAN is never distributed to interfaces outside that VLAN. If hosts in multiple VLANs request the same multicast stream, a separate copy of that multicast stream is distributed to the requesting VLANs.

MVR introduces the concept of a *multicast source VLAN* (MVLAN), which is created by MVR and becomes the only VLAN over which IPTV multicast traffic flows throughout the Layer 2 network. The J-EX Series Switch that is enabled for MVR selectively forward IPTV multicast traffic from interfaces on the MVLAN (source interfaces) to hosts that are connected to interfaces that are not part of the MVLAN. These interfaces are known as *MVR receiver ports*. The MVR receiver ports can receive traffic from a port on the MVLAN but cannot send traffic onto the MVLAN, and they remain in their own VLANs for bandwidth and security reasons.

This topic includes:

How MVR Works

In many ways, MVR is similar to IGMP snooping. Both monitor IGMP join and leave messages and build forwarding tables based on the media access control (MAC) addresses of the hosts sending those IGMP messages. Whereas IGMP snooping operates within a given VLAN to regulate multicast traffic, MVR can operate with hosts on different VLANs in a Layer 2 network to selectively deliver IPTV multicast traffic to requesting hosts, thereby reducing the amount of bandwidth needed to forward multicast traffic.

When you configure an MVLAN, you assign a range of multicast group addresses to it. You then configure other VLANs to be MVR receiver VLANs, which receive multicast streams from the MVLAN. The MVR receiver ports comprise all the interfaces that exist on any of the MVR receiver VLANs. Interfaces that are on the MVLAN itself cannot be MVR receiver ports for that MVLAN.

NOTE: MVR is supported on VLANs running IGMP version 2 (IGMPv2) only.

MVR Modes

MVR operates in two modes: MVR transparent mode and MVR proxy mode. Both modes allow MVR to forward only one copy of a multicast stream to the Layer 2 network.

- MVR Transparent Mode on page 2053
- MVR Proxy Mode on page 2053

MVR Transparent Mode

In MVR transparent mode (the default mode), the switch receives one copy of each IPTV multicast stream and then replicates the stream only to those hosts that want to receive it, while forwarding all other types of multicast traffic without modification. Transparent mode is the default mode.

The switch handles IGMP packets destined for both the multicast source VLAN and multicast receiver VLANs in the same way that it handles them when MVR is not being used. That is, when a host on a VLAN sends IGMP join and leave messages, the switch floods the messages to all router interfaces in the VLAN. Similarly, when a VLAN receives IGMP queries from its router interfaces, it floods the queries to all interfaces in the VLAN.

If a host on a multicast receiver port joins an MVR group on the multicast receiver VLAN, the appropriate bridging entry is added and the MVLAN forwards that group's IPTV multicast traffic on that port (even though that port is not in the MVLAN). Likewise, if a host on a multicast receiver port leaves an MVR group on the multicast receiver VLAN, the appropriate bridging entry is deleted and the MVLAN stops forwarding that group's IPTV multicast traffic on that port. In addition, you can configure the switch to statically install the bridging entries on the multicast receiver VLAN.

MVR Proxy Mode

When you use MVR in proxy mode, the switch acts as a proxy for any MVR group in both the upstream and downstream directions. In the downstream direction, the switch acts as the querier for the groups in the MVR receiver VLANs. In the upstream direction, the switch originates the IGMP reports and leaves and answers IGMP queries from multicast routers. When the MVR receiver VLANs receive IGMP joins and leaves, the switch creates bridging entries on the MVLAN as needed, as it does in MVR transparent mode. In addition, the switch sends out IGMP joins and leaves on the MVLAN based on these bridging entries.

Configuring MVR proxy mode on the MVLAN automatically enables IGMP snooping proxy mode on all MVR receiver VLANs as well as on the MVLAN.

Related Documentation

- Example: Configuring Multicast VLAN Registration on J-EX Series Switches on page 2058
- Configuring Multicast VLAN Registration (CLI Procedure) on page 2068

CHAPTER 76

Examples: IGMP Snooping and Multicast Configuration

Example: Configuring IGMP Snooping on J-EX Series Switches

IGMP snooping regulates multicast traffic in a switched network. With IGMP snooping enabled, a LAN switch monitors the IGMP transmissions between a host (a network device) and a multicast router, keeping track of the multicast groups and associated member ports. The switch uses that information to make intelligent multicast-forwarding decisions and forward traffic to the intended destination interfaces.

Configure IGMP snooping on one or more VLANs to allow the switch to examine IGMP packets and make forwarding decisions based on packet content. By default, IGMP snooping is enabled on J-EX Series switches.

This example describes how to configure IGMP snooping:

Requirements

This example uses the following software and hardware components:

- One J-EX4200-24T switch

Before you configure IGMP snooping, be sure you have:

- Configured the **employee-vlan** VLAN on the switch
- Assigned interfaces **ge-0/0/1**, **ge-0/0/2**, and **ge-0/0/3** to **employee-vlan**

See “Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches” on page 1070.

Overview and Topology

IGMP snooping controls multicast traffic in a switched network. With IGMP snooping enabled, a J-EX Series switch monitors the IGMP transmissions between a host and a multicast router to keep track of the multicast groups and associated member ports. The switch uses this information to make intelligent decisions and forward multicast traffic to the intended destination interfaces.

You can configure IGMP snooping on all interfaces in a VLAN or on individual interfaces. This example shows how to configure IGMP snooping on a J-EX Series switch.

The configuration setup for this example includes the VLAN **employee-vlan** on the switch.

Table 256 on page 2056 shows the components of the topology for this example.

Table 256: Components of the IGMP Snooping Topology

Properties	Settings
Switch hardware	One J-EX4200-24T switch
VLAN name	**employee-vlan**, tag **20**
Interfaces in **employee-vlan**	**ge-0/0/1**, **ge-0/0/2**, **ge-0/0/3**
Multicast IP address for **employee-vlan**	**225.100.100.100**

In this example, the switch is initially configured as follows:

- IGMP snooping is disabled on the VLAN.

Configuration

To configure basic IGMP snooping on a switch:

CLI Quick Configuration

To quickly configure IGMP snooping, copy the following commands and paste them into the switch terminal window:

```
[edit protocols]
set igmp-snooping vlan employee-vlan
set igmp-snooping vlan employee-vlan interface ge-0/0/1 group-limit 50
set igmp-snooping vlan employee-vlan immediate-leave
set igmp-snooping vlan employee-vlan interface ge-0/0/3 static group 225.100.100.100
set igmp-snooping vlan employee-vlan interface ge-0/0/2 multicast-router-interface
set igmp-snooping vlan employee-vlan robust-count 4
```

Step-by-Step Procedure

Configure IGMP snooping:

1. Enable and configure IGMP snooping on the VLAN **employee-vlan**:

   ```
   [edit protocols]
   user@switch# set igmp-snooping vlan employee-vlan
   ```

2. Configure the limit for the number of multicast groups allowed on the **ge-0/0/1** interface to 50.

```
[edit protocols]
user@switch# set igmp-snooping vlan employee-vlan interface ge-0/0/1 group-limit
50
```

3. Configure the switch to immediately remove a group membership from an interface when it receives a leave message from that interface without waiting for any other IGMP messages to be exchanged (IGMPv2 only):

```
[edit protocols]
user@switch# set igmp-snooping vlan employee-vlan immediate-leave
```

4. Statically configure IGMP group membership on a port:

```
[edit protocols]
user@switch# set igmp-snooping vlan employee-vlan interface ge-0/0/3.0 static group
225.100.100.100
```

5. Statically configure an interface as a switching interface toward a multicast router (the interface to receive multicast traffic):

```
[edit protocols]
user@switch# set igmp-snooping vlan employee-vlan interface ge-0/0/2
multicast-router-interface
```

6. Change the number of timeout intervals the switch waits before timing out a multicast group to 4:

```
[edit protocols]
user@switch# set igmp-snooping vlan employee-vlan robust-count 4
```

Results Check the results of the configuration:

```
user@switch# show protocols igmp-snooping
  vlan employee-vlan {
    robust-count 4;
    immediate-leave;
    interface ge-0/0/1 {
      group-limit 50;
    }
    interface ge-0/0/2 {
      multicast-router-interface;
    }
    interface ge-0/0/3 {
      static {
        group 255.100.100.100
      }
    }
  }
```

Related Documentation

- Configuring IGMP Snooping (CLI Procedure) on page 2063
- [edit protocols] Configuration Statement Hierarchy on page 48

Example: Configuring Multicast VLAN Registration on J-EX Series Switches

Multicast VLAN registration (MVR) allows hosts that are not part of a multicast VLAN (MVLAN) to receive multicast streams from the MVLAN, allowing the MVLAN to be shared across the Layer 2 network and eliminating the need to send duplicate multicast streams to each requesting VLAN in the network. Hosts remain in their own VLANs for bandwidth and security reasons.

This example describes how to configure MVR on J-EX Series switches:

Requirements

This example uses the following hardware and software components:

- One J-EX Series switch

Before you configure MVR, be sure you have:

- Configured two or more VLANs on the switch. See “Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches” on page 1070.
- Connected the J-EX Series switch to a network that can transmit IPTV multicast streams from a video server.
- Connected a host that is capable of receiving IPTV multicast streams to an interface in one of the VLANs.

Overview and Topology

In a standard Layer 2 network, a multicast stream received on one VLAN is never distributed to interfaces outside that VLAN. If hosts in multiple VLANs request the same multicast stream, a separate copy of that multicast stream is distributed to the requesting VLANs.

MVR introduces the concept of a *multicast source VLAN* (MVLAN), which is created by MVR and becomes the only VLAN over which multicast traffic flows throughout the Layer 2 network. Multicast traffic can then be selectively forwarded from interfaces on the MVLAN (source ports) to hosts that are connected to interfaces (multicast receiver ports) that are not part of the multicast source VLAN. When you configure an MVLAN, you assign a range of multicast group addresses to it. You then configure other VLANs to be MVR receiver VLANs, which receive multicast streams from the MVLAN. The MVR receiver ports comprise all the interfaces that exist on any of the MVR receiver VLANs.

You can configure MVR to operate in one of two modes: transparent mode (the default mode) or proxy mode. Both modes allow MVR to forward only one copy of a multicast stream to the Layer 2 network.

In transparent mode, the switch receives one copy of each IPTV multicast stream and then replicates the stream only to those hosts that want to receive it, while forwarding all other types of multicast traffic without modification. Figure 1 shows how MVR operates in transparent mode.

In proxy mode, the switch acts as a proxy for the IGMP multicast router in the MVLAN for MVR group memberships established in the MVR receiver VLANs and generates and sends IGMP packets into the MVLAN as needed. Figure 2 shows how MVR operates in proxy mode.

This example shows how to configure MVR in both transparent mode and proxy mode on a J-EX Series switch. The topology includes a video server that is connected to a multicast router, which in turn forwards the IPTV multicast traffic in the MVLAN to the Layer 2 network.

Figure 43 on page 2060 shows the MVR topology in transparent mode. Interfaces P1 and P2 on Switch C belong to service VLAN **s0** and MVLAN **mv0**. Interface P4 of Switch C also belongs to service VLAN **s0**. In the upstream direction of the network, only non-IPTV traffic is being carried in individual customer VLANs of service VLAN **s0**. VLAN **c0** is an example of this type of customer VLAN. IPTV traffic is being carried on MVLAN **mv0**. If any host on any customer VLAN connected to port P4 requests an MVR stream, switch C takes the stream from VLAN **mv0** and replicates that stream onto port P4 with tag **mv0**. IPTV traffic, along with other network traffic, flows form port P4 out to the Digital Subscriber Line Access Multiplexer (DSLAM) **D1**.

Figure 43: MVR Topology in Transparent Mode

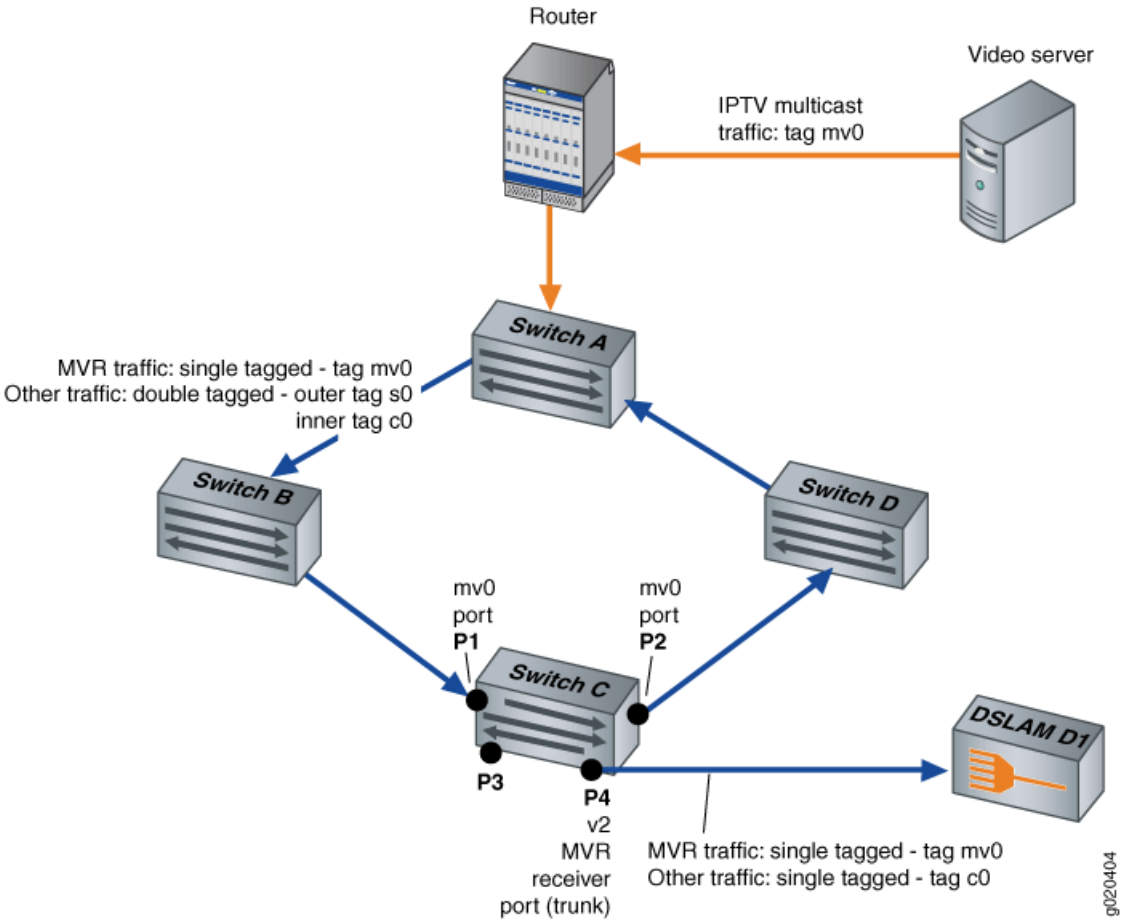

Figure 44 on page 2061 shows the MVR topology in proxy mode. Interfaces P1 and P2 on switch C belong to MVLAN **mv0** and customer VLAN **c0**. Interface P4 on switch C is an access port of customer VLAN **c0**. In the upstream direction of the network, only non-IPTV traffic is being carried on customer VLAN **c0**. Any IPTV traffic requested by hosts on VLAN **c0** is replicated untagged to port P4 based on streams received in MVLAN **mv0**. IPTV traffic flows from port P4 out to an IPTV-enabled device in Host 1. Other traffic, such as data and voice traffic, also flows from port P4 to other network devices in Host 1.

Figure 44: MVR Topology in Proxy Mode

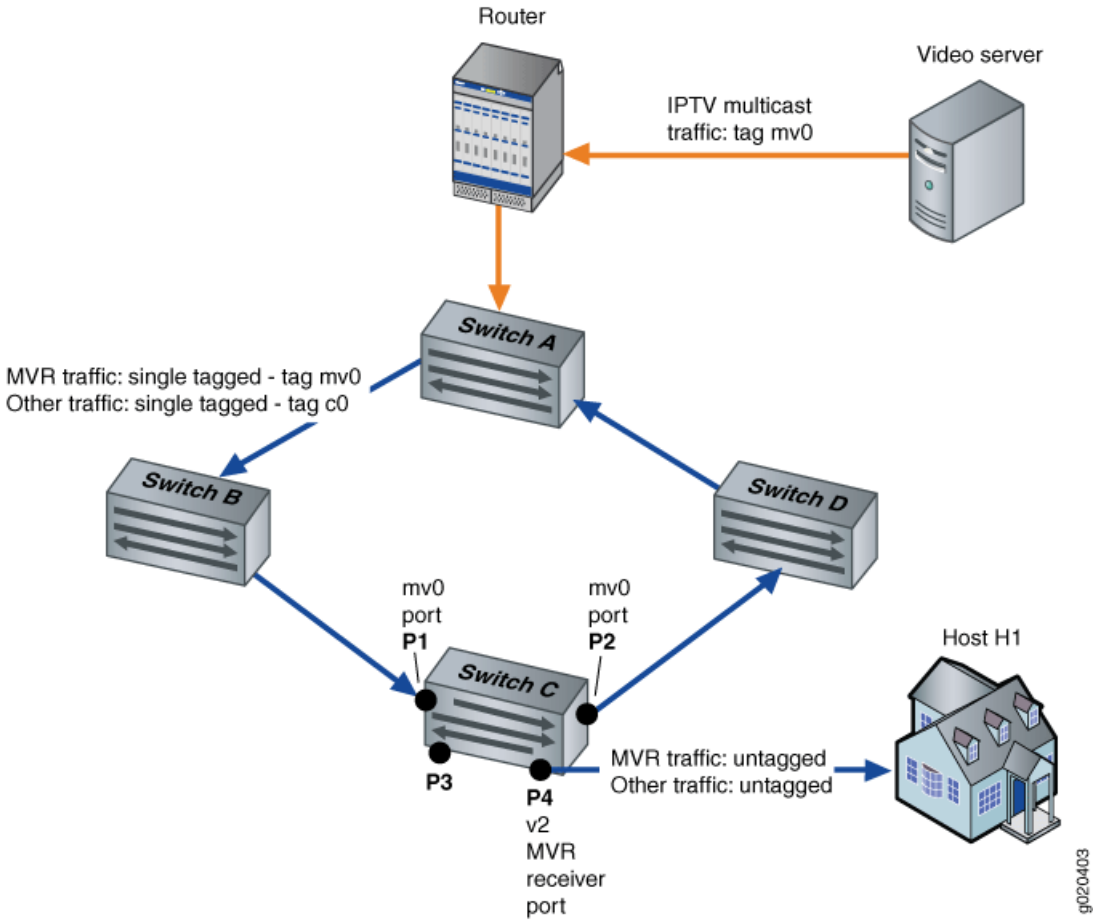

For information on VLAN tagging, see “Understanding Bridging and VLANs on J-EX Series Switches” on page 1041.

Configuration

To configure MVR perform these tasks:

CLI Quick Configuration

To quickly configure MVR in proxy mode, copy the following commands and paste them into the switch terminal window. To quickly configure MVR in transparent mode (the default mode), do not copy and paste the final command line in the following block of lines:

```
[edit protocols igmp-snooping]
set vlan mv0 data-forwarding source groups 225.10.0.0/16
set vlan v2 data-forwarding receiver source-vlans mv0
set vlan v2 data-forwarding receiver install
set vlan mv0 proxy source-address 10.1.1.1
```

Step-by-Step Procedure

To configure MVR, perform these tasks:

1. Configure **mv0** to be an MVLAN:

```
[edit protocols igmp-snooping]
user@switch# set vlan mv0 data-forwarding source groups 225.10.0.0/16
```

2. Configure **v2** to be a multicast receiver VLAN with **mv0** as its source:

```
[edit protocols igmp-snooping]
user@switch# set vlan v2 data-forwarding receiver source-vlans mv0
```

3. (Optional) Install forwarding entries in the multicast receiver VLAN **v2**:

```
[edit protocols igmp-snooping]
user@switch# set vlan v2 data-forwarding receiver install
```

4. (Optional) Configure MVR in proxy mode:

```
[edit protocols igmp-snooping]
user@switch# set vlan mv0 proxy source-address 10.1.1.1
```

Results

Check the results of the configuration:

```
[edit protocols igmp-snooping]
user@switch# show
vlan mv0 {
  proxy {
    source-address 10.1.1.1;
  }
  data-forwarding {
    source {
      groups 225.10.0.0/16;
    }
  }
}
vlan v2 {
  data-forwarding {
    receiver {
      source-vlans mv0;
      install;
    }
  }
}
```

Related Documentation

- Configuring Multicast VLAN Registration (CLI Procedure) on page 2068
- Understanding Multicast VLAN Registration on J-EX Series Switches on page 2052

CHAPTER 77

Configuring IGMP Snooping and Multicast

Configuring IGMP Snooping (CLI Procedure)

IGMP snooping regulates multicast traffic in a switched network. With IGMP snooping enabled, a LAN switch monitors the IGMP transmissions between a host (a network device) and a multicast router, keeping track of the multicast groups and associated member ports. The switch uses that information to make intelligent multicast-forwarding decisions and forward traffic to the intended destination interfaces.

You can configure IGMP snooping on one or more VLANs to allow the switch to examine IGMP packets and make forwarding decisions based on packet content. By default, IGMP snooping is enabled on J-EX Series switches.

NOTE: You cannot configure IGMP snooping on a secondary VLAN.

To enable IGMP snooping and configure individual options as needed for your network by using the CLI:

1. Enable IGMP snooping on a VLAN:

```
[edit protocols]
user@switch# set igmp-snooping vlan employee-vlan
```

2. Configure the limit for the number of multicast groups allowed on the **ge-0/0/1** interface to 50.

```
[edit protocols]
user@switch# set igmp-snooping vlan employee-vlan interface ge-0/0/1 group-limit
50
```

3. Configure the switch to immediately remove a group membership from an interface when it receives a leave message from that interface without waiting for any other IGMP messages to be exchanged (IGMPv2 only):

```
[edit protocols]
user@switch# set igmp-snooping vlan employee-vlan immediate-leave
```

4. Statically configure IGMP group membership on a port:

```
[edit protocols]
user@switch# set igmp-snooping vlan employee-vlan interface ge-0/0/3.0 static group
225.100.100.100
```

5. Statically configure an interface as a switching interface toward a multicast router (the interface to receive multicast traffic):

```
[edit protocols]
user@switch# set igmp-snooping vlan employee-vlan interface ge-0/0/2.0
multicast-router-interface
```

6. Change the number of timeout intervals the switch waits before timing out a multicast group to 4:

```
[edit protocols]
user@switch# set igmp-snooping vlan employee-vlan robust-count 4
```

Related Documentation

- Example: Configuring IGMP Snooping on J-EX Series Switches on page 2055
- Changing the IGMP Snooping Group Query Membership Timeout Value (CLI Procedure) on page 2067
- **show igmp-snooping membership on page 2181**
- **show igmp-snooping route on page 2183**
- **show igmp-snooping statistics on page 2185**
- **show igmp-snooping vlans on page 2187**
- IGMP Snooping on J-EX Series Switches Overview on page 2047

Configuring IGMP Snooping (J-Web Procedure)

IGMP snooping regulates multicast traffic in a switched network. With IGMP snooping enabled, the J-EX Series switch monitors the IGMP transmissions between a host (a network device) and a multicast router, keeping track of the multicast groups and associated member interfaces. The switch uses that information to make intelligent multicast-forwarding decisions and forward traffic to the intended destination interfaces.

You can configure IGMP snooping on one or more VLANs to allow the switch to examine IGMP packets and make forwarding decisions based on packet content. By default, IGMP snooping is enabled on J-EX Series switches.

To enable IGMP snooping and configure individual options using the J-Web interface:

1. Select **Configure** > **Switching** > **IGMP Snooping**.

NOTE: After you make changes to the configuration in this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See "Using the Commit Options to Commit Configuration Changes (J-Web Procedure)" on page 334 for details about all commit options.

2. Click one:
   - **Add**—Creates an IGMP snooping configuration for the VLAN.
   - **Edit**—Modifies an IGMP snooping configuration for the VLAN.
   - **Delete**—Deletes a selected VLAN from the IGMP snooping configuration.

   When you are adding or editing an IGMP snooping configuration, enter information as described in Table 257 on page 2065
3. Click **OK** to apply changes to the configuration or click **Cancel** to cancel without saving changes.

To disable IGMP snooping on a VLAN, select the VLAN from the list and click **Disable**.

Table 257: IGMP Snooping Configuration Fields

Field	Function	Your Action
VLAN Name	Specifies the VLAN on which to enable IGMP snooping.	Select a VLAN from the list to add it to the snooping configuration.
Immediate Leave	Immediately removes a multicast group membership from an interface when it receives a leave message from that interface without waiting for any other IGMP messages to be exchanged (IGMPv2 only).	To enable the option, select the check box. To disable the option, clear the check box.
Robust Count	Specifies the number of timeout intervals the switch waits before timing out a multicast group.	Type a value.

Table 257: IGMP Snooping Configuration Fields *(continued)*

Field	Function	Your Action
Interfaces List	Statically configures an interface as a switching interface toward a multicast router (the interface to receive multicast traffic).	Click one: • Add—Adds an interface to the IGMP snooping configuration. 1. Select an interface from the list. 2. Select **Multicast Router Interface**. 3. Type the maximum number of groups an interface can join. 4. In **Static**, choose one: ◦ Click **Add**, type a group IP address, and click **OK**. ◦ Select a group and click **Remove** to remove the group membership. • Edit—Edits the interface settings for the IGMP snooping configuration. • Remove—Deletes an interface configured for IGMP snooping.

Related Documentation

- Example: Configuring IGMP Snooping on J-EX Series Switches on page 2055
- Configuring IGMP Snooping (CLI Procedure) on page 2063
- Changing the IGMP Snooping Group Query Membership Timeout Value (CLI Procedure) on page 2067
- IGMP Snooping on J-EX Series Switches Overview on page 2047

Changing the IGMP Snooping Group Query Membership Timeout Value (CLI Procedure)

Generally, you do not need to explicitly set the group membership timeout value for IGMP snooping groups on a J-EX Series switch. The group membership timeout value, which determines how long the switch waits before removing an IGMP snooping group from its multicast cache table. is implicitly set to 260 seconds when you configure IGMP snooping.

When you enable IGMP snooping on a switch, the **query-interval** and **query-response-interval** values are set to their default values and are applied to all VLANs created on the switch. The default values are:

- **query-interval**—125 seconds
- **query-response-interval**—10 seconds

The software automatically calculates the group membership timeout value for an IGMP snooping-enabled switch by multiplying the **query-interval** value by 2 and then adding the **query-response-interval** value. For example, using the default values: (125 x 2) + 10) = 260.

If you need to explicitly set the group membership timeout value, you reset the **query-interval** and **query-response-interval** values at the **[edit protocols igmp]** hierarchy level. (Notice that you are not resetting the values at the **[edit protocols igmp-snooping]** hierarchy level.) When you reset these values, the IGMP snooping configuration inherits the new values and recalculates the group membership timeout value accordingly. For more information on changing these values, see the *Junos OS Multicast Protocols Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.

To change the IGMP snooping group membership timeout value to 350:

1. Configure the **query-interval** value to be 150:

```
[edit protocols]
user@switch# set igmp query-interval 150
```

2. Configure the **query-response-interval** value to be 50:

```
[edit protocols]
user@switch# set igmp query-response-interval 50
```

Related Documentation

- Example: Configuring IGMP Snooping on J-EX Series Switches on page 2055
- Verifying That the IGMP Snooping Group Query Timeout Value Has Been Changed Correctly on page 2070
- Configuring IGMP Snooping (CLI Procedure) on page 2063
- Configuring IGMP Snooping (J-Web Procedure) on page 2064

Configuring Multicast VLAN Registration (CLI Procedure)

Multicast VLAN registration (MVR) allows hosts that are not part of a multicast source VLAN (MVLAN) to still receive multicast streams from the MVLAN, allowing an MVLAN to be shared across a Layer 2 network. Hosts remain in their own VLANs for bandwidth and security reasons but are able to receive multicast streams from the MVLAN.

You can configure one or more VLANs on a switch to be MVLANs or MVR receiver VLANs. By default, MVR is not configured on J-EX Series switches.

NOTE: MVR is supported on VLANs running IGMP version 2 (IGMPv2) only.

NOTE: When configuring MVR, the following restrictions apply:

- You cannot enable multicast protocols on VLAN interfaces that are members of MVLANs.
- If you configure an MVLAN in proxy mode, IGMP snooping proxy mode will be automatically enabled on all MVR receiver VLANs of this MVLAN. If a VLAN is an MVR receiver VLAN for multiple MVLANs, all of the MVLANs must have proxy mode enabled or all must have proxy mode disabled. You can enable proxy mode only on VLANs that are configured as MVR source VLANs and that are not configured for Q-in-Q tunneling.
- After you configure a VLAN as an MVLAN, that VLAN is no longer available for other uses.

To configure MVR:

1. Configure the VLAN named **mv0** to be an MVLAN:

```
[edit protocols]
user@switch# set igmp-snooping vlan mv0 data-forwarding source groups 225.10.0.0/16
```

2. Configure the MVLAN **mv0** to be a proxy VLAN:

```
[edit protocols]
user@switch# set igmp-snooping vlan mv0 proxy source–address 10.0.0.1
```

3. Configure the VLAN named **v2** to be an MVR receiver VLAN:

```
[edit protocols]
user@switch# set igmp-snooping vlan v2 data-forwarding receiver source-vlans mv0
```

4. Install forwarding entries in the MVR receiver VLAN:

```
[edit protocols]
user@switch# set igmp-snooping vlan mv0 data-forwarding receiver install
```

Related Documentation

- Example: Configuring Multicast VLAN Registration on J-EX Series Switches on page 2058
- Understanding Multicast VLAN Registration on J-EX Series Switches on page 2052

CHAPTER 78

Verifying IGMP Snooping and Multicast

Monitoring IGMP Snooping

Purpose Use the monitoring feature to view status and information about IGMP snooping configuration on your J-EX Series switch.

Action To display IGMP snooping details in the J-Web interface, select **Monitor > Switching > IGMP Snooping**.

To display IGMP snooping details in the CLI, enter the following commands:

- **show igmp-snooping vlans**
- **show igmp-snooping statistics**
- **show igmp-snooping route**

Meaning Table 258 on page 2069 summarizes the IGMP snooping details displayed.

Table 258: Summary of IGMP Snooping Output Fields

Field	Values
IGMP Snooping Monitor	
VLAN	The VLAN for which IGMP snooping is enabled.
Interfaces	Indicates the interfaces configured as switching interfaces that are associated with the multicast router.
Groups	Indicates the number of the multicast groups learned by the VLAN.
MRouters	Specifies the multicast router.
Receivers	Specifies the multicast receiver.
IGMP Route Information	

Table 258: Summary of IGMP Snooping Output Fields *(continued)*

Field	Values
VLAN	The VLAN for which IGMP snooping is enabled.
Group	Indicates the multicast groups learned by the VLAN.
Next-Hop	Specifies the next hop assigned by the switch after performing the route lookup.

Related Documentation

- **show igmp-snooping vlans on page 2187**
- **show igmp-snooping statistics on page 2185**
- **show igmp-snooping route on page 2183**
- Configuring IGMP Snooping (CLI Procedure) on page 2063
- Example: Configuring IGMP Snooping on J-EX Series Switches on page 2055

Verifying That the IGMP Snooping Group Query Timeout Value Has Been Changed Correctly

Purpose Verify that the IGMP snooping group query timeout value has been changed correctly from its default value.

Action Display the IGMP protocol information:

```
user@switch> show configuration protocols igmp
query-interval 150;
query-response-interval 50;
accounting;
interface vlan.43 {
    version 2;
}
```

Display the IGMP snooping membership information, which contains the group query timeout value that was derived from the IGMP configuration:

```
user@switch> show show igmp-snooping membership detail
VLAN: v43 Tag: 43 (Index: 4)
      Group: 225.0.0.1
      Receiver count: 1, Flags: <v2-hosts>
          ge-0/0/15.0 Uptime: 00:00:05 timeout: 350
```

Meaning When you enable IGMP snooping on a switch, the **query-interval** and **query-response-interval** values are set to their default values and are applied to all VLANs created on the switch. The IGMP snooping group timeout value is derived from these default settings. Based on the default values, the initial IGMP snooping group query timeout value is 260.

To change the group query timeout value, change the **query-interval** and **query-response-interval** values at the **[edit protocols igmp]** hierarchy level. The IGMP snooping group query timeout value is then recalculated based on the new IGMP configuration settings.

The output from the **show protocols igmp** command shows the revised IGMP configuration settings for **query-interval** and **query-response-interval**. You know that these values have been revised because they are different from the default values. The output from the **show igmp-snooping membership detail** command shows the revised group query timeout value, **350**, which was derived from the new IGMP configuration settings.

Related Documentation

- Changing the IGMP Snooping Group Query Membership Timeout Value (CLI Procedure) on page 2067

CHAPTER 79

Configuration Statements for IGMP Snooping and Multicast

[edit protocols] Configuration Statement Hierarchy

```
protocols {
  connections {
    remote-interface-switch connection-name {
      interface interface-name.unit-number;
      transmit-lsp label-switched-path;
      receive-lsp label-switched-path;
    }
  }
  dot1x {
    authenticator {
      authentication-profile-name profile-name;
      interface (all | [ interface-names ]) {
        disable;
        guest-vlan ( vlan-id | vlan-name);
        mac-radius <restrict>;
        maximum-requests number;
        no-reauthentication;
        quiet-period seconds;
        reauthentication {
          interval seconds;
        }
        retries number;
        server-fail (deny | permit | use-cache | vlan-id | vlan-name);
        server-reject-vlan (vlan-id | vlan-name);
        server-timeout seconds;
        supplicant (multiple | single | single-secure);
        supplicant-timeout seconds;
        transmit-period seconds;
      }
      static mac-address {
        interface interface-name;
        vlan-assignment (vlan-id |vlan-name);
      }
  }
  gvrp {
```

```
    <enable | disable>;
    interface (all | [interface-name]) {
      disable;
    }
    join-timer millseconds;
    leave-timer milliseconds;
    leaveall-timer milliseconds;
  }
  igmp-snooping {
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>
        <match regex>;
      flag flag (detail | disable | receive | send);
    }
    vlan (vlan-id | vlan-number) {
      data-forwarding {
        source {
          groups group-prefix;
        }
        receiver {
          source-vlans vlan-list;
          install ;
        }
      }
      disable {
        interface interface-name
      }
      immediate-leave;
      interface interface-name {
        group-limit limit;
        multicast-router-interface;
        static {
          group ip-address;
        }
      }
      proxy ;
      query-interval seconds;
      query-last-member-interval seconds;
      query-response-interval seconds;
      robust-count number;
    }
  }
  lldp {
    disable;
    advertisement-interval seconds;
    hold-multiplier number;
    interface (all | interface-name) {
      disable;
    }
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>
        <match regex>;
      flag flag (detail | disable | receive | send);
    }
  }
  lldp-med {
```

```
    disable;
    fast-start number;
    interface (all | interface-name) {
      disable;
      location {
        elin number;
        civic-based {
          what number;
          country-code code;
          ca-type {
            number {
              ca-value value;
            }
          }
        }
      }
    }
  }
  mpls {
    interface ( all | interface-name );
    label-switched-path lsp-name to remote-provider-edge-switch;
    path destination {
      <address | hostname> <strict | loose>
    }
  mstp {
    disable;
    bpdu-block-on-edge;
    bridge-priority priority;
    configuration-name name;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
      disable;
      bpdu-timeout-action {
        block;
        alarm;
      }
      cost  cost;
      edge;
      mode mode;
      no-root-port;
      priority priority;
    }
    max-age seconds;
    max-hops hops;
    msti msti-id {
      vlan (vlan-id | vlan-name);
      interface interface-name {
        disable;
        cost  cost;
        edge;
        mode mode;
        priority priority;
      }
    }
    revision-level revision-level;
```

```
    traceoptions {
      file filename <files number > <size size> <no-stamp | world-readable |
        no-world-readable>;
      flag flag;
    }
}
mvrp {
  disable
  interface (all | interface-name) {
    disable;
    join-timer milliseconds;
    leave-timer milliseconds;
    leaveall-timer milliseconds;
    registration (forbidden | normal);
  }
  no-dynamic-vlan;
  traceoptions {
    file filename <files number > <size size> <no-stamp | world-readable |
      no-world-readable>;
    flag flag;
  }
}
oam {
  ethernet{
    connectivity-fault-management {
      action-profile profile-name {
        default-actions {
          interface-down;
        }
      }
      linktrace {
        age (30m | 10m | 1m | 30s | 10s);
        path-database-size path-database-size;
      }
      maintenance-domain domain-name {
        level number;
        mip-half-function (none | default |explicit);
        name-format (character-string | none | dns | mac+2oct);
        maintenance-association ma-name {
          continuity-check {
            hold-interval minutes;
            interval (10m | 10s | 1m | 1s| 100ms);
            loss-threshold number;
          }
          mep mep-id {
            auto-discovery;
            direction down;
            interface interface-name;
            remote-mep mep-id {
              action-profile profile-name;
            }
          }
        }
      }
    }
    link-fault-management {
```

```
        action-profile profile-name;
          action {
            syslog;
            link-down;
          }
          event {
            link-adjacency-loss;
            link-event-rate;
            frame-error count;
            frame-period count;
            frame-period-summary count;
            symbol-period count;
          }
          interface interface-name {
            link-discovery (active | passive);
            pdu-interval interval;
            event-thresholds threshold-value;
            remote-loopback;
            event-thresholds {
              frame-errorcount;
              frame-period count;
              frame-period-summary count;
              symbol-period count;
            }
          }
          negotiation-options {
            allow-remote-loopback;
            no-allow-link-events;
          }
      }
    }
  }
  rstp {
    disable;
    bpdu-block-on-edge;
    bridge-priority priority;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
      disable;
      bpdu-timeout-action {
        block;
        alarm;
      }
      cost  cost;
      edge;
      mode mode;
      no-root-port;
      priority priority;
    }
    max-age seconds;
    }
    traceoptions {
      file filename <files number > <size size> <no-stamp | world-readable |
        no-world-readable>;
      flag flag;
```

```
    }
  }
  sflow {
    agent-id
    collector {
      ip-address;
      udp-port port-number;
    }
    disable;
    interfaces interface-name {
      disable;
      polling-interval seconds;
      sample-rate number;
    }
    polling-interval seconds;
    sample-rate number;
    source-ip
  }
  stp {
    disable;
    bridge-priority priority;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
      disable;
      bpdu-timeout-action {
        block;
        alarm;
      }
      cost  cost;
      edge;
      mode mode;
      no-root-port;
      priority priority;
    }
    max-age seconds;
  }
    traceoptions {
      file filename <files number > <size size> <no-stamp | world-readable |
        no-world-readable>;
      flag flag;
    }
  vstp {
    bpdu-block-on-edge;
    disable;
    force-version stp;
    vlan (all | vlan-id | vlan-name) {
      bridge-priority priority;
      forward-delay seconds;
      hello-time seconds;
      interface (all | interface-name) {
        bpdu-timeout-action {
          alarm;
          block;
        }
        cost  cost;
```

```
            disable;
            edge;
            mode mode;
            no-root-port;
            priority priority;
          }
          max-age seconds;
          traceoptions {
            file filename <files number > <size size> <no-stamp | world-readable |
              no-world-readable>;
            flag flag;
          }
        }
      }
    }
```

Related Documentation

- 802.1X for J-EX Series Switches Overview on page 2253
- Example: Configure Automatic VLAN Administration Using GVRP on page 1087
- Understanding MAC RADIUS Authentication on J-EX Series Switches
- Understanding Server Fail Fallback and 802.1X Authentication on J-EX Series Switches on page 2258
- IGMP Snooping on J-EX Series Switches Overview on page 2047
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261
- Understanding MSTP for J-EX Series Switches on page 1277
- Understanding Multiple VLAN Registration Protocol (MVRP) on J-EX Series Switches on page 1054
- Understanding Ethernet OAM Connectivity Fault Management for a J-EX Series Switch on page 3463
- Understanding Ethernet OAM Link Fault Management for a J-EX Series Switch on page 3427
- Understanding RSTP for J-EX Series Switches on page 1276
- Understanding STP for J-EX Series Switches on page 1275
- Understanding How to Use sFlow Technology for Network Monitoring on a J-EX Series Switch on page 3283
- Understanding VSTP for J-EX Series Switches on page 1281

accounting (Per Interface)

Syntax (accounting | no-accounting);

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp **interface** *interface-name*],
[edit protocols igmp **interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable or disable the collection of IGMP join and leave event statistics for an interface.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Recording IGMP Join and Leave Events

accounting (Protocol)

Syntax accounting;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp],
[edit protocols igmp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable the collection of IGMP join and leave event statistics on the system.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Recording IGMP Join and Leave Events

address (Anycast RPs)

Syntax address *address* <forward-msdp-sa>;

Hierarchy Level [edit logical-systems *logical-system-name* protocols **pim rp local** (inet | inet6) **anycast-pim rp-set**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim rp local** (inet | inet6) **anycast-pim rp-set**],
[edit protocols **pim rp local** (inet | inet6) **anycast-pim rp-set**],
[edit routing-instances *routing-instance-name* protocols **pim rp local** (inet | inet6) **anycast-pim rp-set**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the anycast rendezvous point (RP) addresses in the RP set. Multiple addresses can be configured in an RP set. If the RP has peer Multicast Source Discovery Protocol (MSDP) connections, then the RP must forward MSDP source active (SA) messages.

Options ***address***—RP address in an RP set.

forward-msdp-sa—(Optional) Forward MSDP SAs to this address.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

address (Local RPs)

Syntax address *address*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols **pim rp local family** (inet | inet6)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim rp local family** (inet | inet6)],
[edit protocols **pim rp local family** (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols **pim rp local family** (inet | inet6)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the local rendezvous point (RP) address.

Options ***address***—Local RP address.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Local PIM RPs

anycast-pim

Syntax

```
anycast-pim {
  rp-set {
    address address <forward-msdp-sa>;
  }
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols **pim rp local family** (inet | inet6)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim rp local family** (inet | inet6)],
[edit protocols **pim rp local family** (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols **pim rp local family** (inet | inet6)]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure properties for anycast RP using PIM.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring PIM Anycast with MSDP

assert-timeout

Syntax assert-timeout *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols **pim**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim**],
[edit protocols **pim**],
[edit routing-instances *routing-instance-name* protocols **pim**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Multicast routing devices running PIM sparse mode often forward the same stream of multicast packets onto the same LAN through the rendezvous-point tree (RPT) and shortest-path tree (SPT). PIM assert messages help routing devices determine which routing device forwards the traffic and prunes the RPT for this group. By default, routing devices enter an assert cycle every 180 seconds. You can configure this assert timeout to be between 5 and 210 seconds.

Options ***seconds***—Time for routing device to wait before another assert message cycle.

Range: 5 through 210 seconds

Default: 180 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring the PIM Assert Timeout

auto-rp

Syntax

```
auto-rp {
  (announce | discovery | mapping);
  (mapping-agent-election | no-mapping-agent-election);
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols **pim rp**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim rp**],
[edit protocols **pim rp**],
[edit routing-instances *routing-instance-name* protocols **pim rp**]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure automatic RP announcement and discovery.

Options

announce—Configures the routing device to listen only for mapping packets and also to advertise itself if it is an RP.

discovery—Configures the routing device to listen only for mapping packets.

mapping—Configures the routing device to announce, listens for and generates mapping packets, and announces that the routing device is eligible to be an RP.

The remaining statement is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM Auto-RP

bootstrap

Syntax

```
bootstrap {
  family (inet | inet6) {
    export [ policy-names ];
    import [ policy-names ];
    priority number;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim rp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp],
[edit protocols pim rp],
[edit routing-instances *routing-instance-name* protocols pim rp]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure parameters to control bootstrap routers and messages.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM Bootstrap Properties

bootstrap-export

Syntax

```
bootstrap-export [ policy-names ];
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols pim rp],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp],
[edit protocols pim rp],
[edit routing-instances *routing-instance-name* protocols pim rp]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Apply one or more export policies to control outgoing PIM bootstrap messages.

Options

policy-names—Name of one or more import policies.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM Bootstrap Properties
- **bootstrap-import on page 2086**

bootstrap-import

Syntax bootstrap-import [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols **pim rp**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim rp**],
[edit protocols **pim rp**],
[edit routing-instances *routing-instance-name* protocols **pim rp**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply one or more import policies to control incoming PIM bootstrap messages.

Options ***policy-names***—Name of one or more import policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring PIM Bootstrap Properties
- **bootstrap-export on page 2085**

bootstrap-priority

Syntax bootstrap-priority *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols **pim rp**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim rp**],
[edit protocols **pim rp**],
[edit routing-instances *routing-instance-name* protocols **pim rp**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure whether this routing device is eligible to be a bootstrap router. In the case of a tie, the routing device with the highest IP address is elected to be the bootstrap router.

Options ***number***—Priority for becoming the bootstrap router. A value of 0 means that the routing device is not eligible to be the bootstrap router.

Range: 0 through 255

Default: 0

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Junos OS Multicast Protocols Configuration Guide*

data-forwarding

Syntax

```
data-forwarding {
  source {
    groups group-prefix;
  }
  receiver {
    source-vlans vlan-list;
    install;
  }
}
```

Hierarchy Level [edit protocols **igmp-snooping vlan** *vlan-id* | *vlan-number*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the VLAN to be a multicast source VLAN (MVLAN) or a multicast VLAN registration (MVR) receiver VLAN. Each data-forwarding VLAN, which can be a multicast source VLAN (MVLAN) or a multicast receiver VLAN, must have exactly one source statement or exactly one receiver statement. A data-forwarding VLAN can operate only in IGMPv2 mode.

The remaining statements are explained separately.

Default Disabled.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [edit protocols] Configuration Statement Hierarchy on page 48
- Example: Configuring Multicast VLAN Registration on J-EX Series Switches on page 2058
- Configuring Multicast VLAN Registration (CLI Procedure) on page 2068

dense-groups

Syntax
```
dense-groups {
    addresses;
}
```

Hierarchy Level
[edit logical-systems *logical-system-name* protocols **pim**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim**],
[edit protocols **pim**],
[edit routing-instances *routing-instance-name* protocols **pim**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure which groups are operating in dense mode.

Options ***addresses***—Address of groups operating in dense mode.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring PIM Sparse-Dense Mode Properties

disable

Syntax
```
disable {
    interface interface-name
}
```

Hierarchy Level [edit protocols **igmp-snooping vlan** *vlan-id* | *vlan-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable IGMP snooping on all interfaces in a VLAN or on a specific VLAN interface.

Default If you do not specify an interface, all interfaces in the given VLAN are disabled.

Options ***interface-name*** —Name of the interface.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring IGMP Snooping on J-EX Series Switches on page 2055
- Configuring IGMP Snooping (CLI Procedure) on page 2063

disable (PIM)

Syntax disable;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols **pim**],
[edit logical-systems *logical-system-name* protocols **pim** family (inet | inet6)],
[edit logical-systems *logical-system-name* protocols **pim interface** *interface-name*],
[edit logical-systems *logical-system-name* protocols **pim rp local** family (inet | inet6)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim rp local** family (inet | inet6)],
[edit protocols **pim**],
[edit protocols **pim** family (inet | inet6)],
[edit protocols **pim interface** *interface-name*],
[edit protocols **pim rp local** family (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols **pim**],
[edit routing-instances *routing-instance-name* protocols **pim** family (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols **pim interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols **pim rp local** family (inet | inet6)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Explicitly disable PIM at the protocol, interface or family hierarchy levels.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Disabling PIM
- family (Disable PIM)

disable

Syntax disable;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols igmp **interface** *interface-name*],
[edit protocols igmp **interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable IGMP on the system.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Disabling IGMP

dr-election-on-p2p

Syntax dr-election-on-p2p;

Hierarchy Level [edit logical-systems *logical-system-name* protocols **pim**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim**],
[edit protocols **pim**],
[edit routing-instances *routing-instance-name* protocols **pim**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable PIM designated router (DR) election on point-to-point (P2P) links.

Default No PIM DR election is performed on P2P links.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring PIM Designated Router Election on Point-to-Point Links

dr-register-policy

Syntax dr-register-policy [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols **pim rp**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim rp**],
[edit protocols **pim rp**],
[edit routing-instances *routing-instance-name* protocols **pim rp**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply one or more policies to control outgoing PIM register messages.

Options ***policy-names***—Name of one or more import policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Register Message Filtering on a PIM RP or DR
- **rp-register-policy on page 2125**

embedded-rp

Syntax

```
embedded-rp {
  group-ranges {
    destination-ip-prefix</prefix-length>;
  }
  maximum-rps limit;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols **pim rp**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim rp**],
[edit protocols **pim rp**],
[edit routing-instances *routing-instance-name* protocols **pim rp**]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure properties for embedded IP version 6 (IPv6) RPs.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM Embedded RP for IPv6

export (Bootstrap)

Syntax

```
export [ policy-names ];
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols **pim rp bootstrap family** (inet | inet6)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim rp bootstrap family** (inet | inet6)],
[edit protocols **pim rp bootstrap family** (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols **pim rp bootstrap family** (inet | inet6)]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Apply one or more export policies to control outgoing PIM bootstrap messages.

Options

policy-names—Name of one or more import policies.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM Bootstrap Properties
- **import (Bootstrap) on page 2103**

family (Bootstrap)

Syntax

```
family (inet | inet6) {
    export [ policy-names ];
    number;
    [ policy-names ];
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols **pim rp bootstrap**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim rp bootstrap**],
[edit protocols **pim rp bootstrap**],
[edit routing-instances *routing-instance-name* protocols **pim rp bootstrap**]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure which IP protocol type bootstrap properties to apply.

Options

inet—Apply IP version 4 (IPv4) local RP properties.

inet6—Apply IPv6 local RP properties.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM Bootstrap Properties

family (Local RP)

Syntax

```
family (inet | inet6) {
  disable;
  address address;
  anycast-pim {
    local-address address;
    rp-set {
      address address <forward-msdp-sa>;
    }
  }
  group-ranges {
    destination-ip-prefix</prefix-length>;
  }
  hold-time seconds;
  priority number;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols **pim rp local**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim rp local**],
[edit protocols **pim rp local**],
[edit routing-instances *routing-instance-name* protocols **pim rp local**]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure which IP protocol type local RP properties to apply.

Options

inet—Apply IP version 4 (IPv4) local RP properties.

inet6—Apply IPv6 local RP properties.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Junos OS Multicast Protocols Configuration Guide*

graceful-restart

Syntax

```
graceful-restart {
  disable;
  restart-duration seconds;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols **pim**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim**],
[edit protocols **pim**],
[edit routing-instances *routing-instance-name* protocols **pim**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure PIM sparse mode graceful restart.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM Sparse Mode Graceful Restart

group

Syntax

group *ip-address*;

Hierarchy Level

[edit protocols **igmp-snooping vlan** *vlan-id* | *vlan-name* **interface** *interface-name* **static**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a static multicast group using a valid IP multicast address.

Default None.

Options ***ip-address*** —IP address of the multicast group receiving data on an interface.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring IGMP Snooping on J-EX Series Switches on page 2055
- Configuring IGMP Snooping (CLI Procedure) on page 2063

group

Syntax

```
group multicast-group-address {
  exclude;
  group-count number;
  group-increment increment;
  source ip-address {
    source-count number;
    source-increment increment;
  }
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp **interface** *interface-name* **static**], [edit protocols igmp **interface** *interface-name* **static**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the IGMP multicast group address and (optionally) the source address for the multicast group being statically configured on an interface.

NOTE: You must specify a unique address for each group.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Enabling IGMP Static Group Membership

group-limit

Syntax group-limit *limit*;

Hierarchy Level [edit protocols **igmp-snooping vlan** *vlan-id* | *vlan-number* **interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a limit for the number of multicast groups allowed on the specified interface. After this limit is reached, new reports are ignored and related flows are not flooded on the interface.

Default No group limits are configured.

Options ***limit*** —Number that represents the maximum number of multicast groups allowed on the specified interface.

Range: 0 through 65535

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring IGMP Snooping on J-EX Series Switches on page 2055
- Configuring IGMP Snooping (CLI Procedure) on page 2063
- Configuring IGMP Snooping (J-Web Procedure) on page 2064
- **group on page 2094**

group-ranges

Syntax

```
group-ranges {
  destination-ip-prefix</prefix-length>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols **pim rp embedded-rp**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim rp embedded-rp**],
[edit protocols **pim rp embedded-rp**],
[edit protocols **pim rp local family** (inet | inet6)],
[edit protocols **pim rp static** address *address*],
[edit routing-instances *routing-instance-name* protocols **pim rp embedded-rp**],
[edit routing-instances *routing-instance-name* protocols **pim rp local family** (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols **pim rp static** address *address*]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure the address ranges of the multicast groups for which this routing device can be an RP.

Default

The routing device is eligible to be the RP for all IPv4 or IPv6 groups (**224.0.0.0/4** or **FF70::/12** to **FFF0::/12**).

Options

destination-mask—Addresses or address ranges for which this routing device can be an RP.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Local PIM RPs
- Configuring PIM Embedded RP for IPv6

groups

Syntax groups *group-prefix*;

Hierarchy Level [edit protocols **igmp-snooping vlan** *vlan-id* | *vlan-number* **data-forwarding source**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the IP address range of the multicast VLAN (MVLAN) source interfaces.

Default Disabled.

Options ***group-prefix***—IP address range of the source group. Each MVLAN must have exactly one **groups** statement. If there are multiple MVLANs on the switch, their group ranges must be unique.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [edit protocols] Configuration Statement Hierarchy on page 48
- Example: Configuring Multicast VLAN Registration on J-EX Series Switches on page 2058
- Configuring Multicast VLAN Registration (CLI Procedure) on page 2068

hello-interval

Syntax hello-interval *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols **pim interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim interface** *interface-name*],
[edit protocols **pim interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols **pim interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify how often the router sends PIM hello packets out of an interface.

Options ***seconds***—Length of time between PIM hello packets.

Range: 0 through 255

Default: 30 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Modifying the PIM Hello Interval
- **hold-time on page 2099**

hold-time

Syntax hold-time *seconds*;

Hierarchy Level [edit protocols **pim rp local family** (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols **pim rp local family** (inet | inet6)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the time period for which a neighbor is to consider the sending routing device (this routing device) to be operative (up).

Options *seconds*—Hold time.

Range: 0 through 255

Default: 0 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Local PIM RPs

igmp-snooping

Syntax

```
igmp-snooping {
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable> <match
      regex>;
    flag flag (detail | disable | receive | send);
  }
  vlan vlan-id | vlan-name {
    data-forwarding {
      source {
        groups group-prefix;
      }
      receiver {
        source-vlans vlan-list;
        install ;
      }
    }
    disable {
      interface interface-name;
    }
    immediate-leave;
    interface interface-name {
      group-limit limit;
      multicast-router-interface;
      static {
        group ip-address;
      }
    }
    proxy ;
    query-interval seconds;
    query-last-member-interval seconds;
    query-response-interval seconds;
    robust-count number;
  }
}
```

Hierarchy Level [edit protocols]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable and configure IGMP snooping on J-EX Series switches.

The remaining statements are explained separately.

Default IGMP snooping is enabled by default.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring IGMP Snooping on J-EX Series Switches on page 2055
- Configuring IGMP Snooping (CLI Procedure) on page 2063

immediate-leave

Syntax immediate-leave;

Hierarchy Level [edit protocols **igmp-snooping vlan** *vlan-id* | *vlan-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description (Applies only to switches running IGMPv2.) After the switch receives a leave group membership message from a host, immediately remove the group membership from the interface without waiting for any other IGMP messages to be exchanged.

NOTE: When configuring this statement, ensure that the IGMP interface has only one IGMP host connected. If more than one IGMPv2 host is connected to the switch through the same interface and one of the hosts sends a leave message, the switch removes all hosts on the interface from the multicast group. The switch loses contact with the hosts in the multicast group that did not send a leave message until they send join requests in response to the next general multicast listener query from the router.

Default The immediate-leave feature is disabled.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring IGMP Snooping on J-EX Series Switches on page 2055
- Configuring IGMP Snooping (CLI Procedure) on page 2063

immediate-leave

Syntax immediate-leave;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp **interface** *interface-name*], [edit protocols igmp **interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description When this statement is enabled on a routing device running IGMP version 2 (IGMPv2), after the routing device receives a leave group membership message from a host associated with the interface, the routing device immediately removes the group membership from the interface and suppresses the sending of any group-specific queries for the multicast group.

When this statement is enabled on a routing device running IGMP version 3 (IGMPv3), after the routing device receives a report with the type BLOCK_OLD_SOURCES, the routing device suppresses the sending of group-and-source queries but relies on the Junos OS-supported host tracking mechanism to determine whether or not it removes a particular source group membership from the interface.

NOTE: When issuing this command on IGMPv2 interfaces, ensure that the IGMP interface has only one IGMP host connected. If more than one IGMPv2 host is connected to a LAN through the same interface, and one host sends a done message, the routing device removes all hosts on the interface from the multicast group. The routing device loses contact with the hosts that properly remain in the multicast group until they send join requests in response to the next general multicast listener query from the router.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Specifying Immediate-Leave Host Removal for IGMP

import (Bootstrap)

Syntax import [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols **pim rp bootstrap** (inet | inet6)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim rp bootstrap** (inet | inet6)],
[edit protocols **pim rp bootstrap** (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols **pim rp bootstrap** (inet | inet6)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply one or more import policies to control incoming PIM bootstrap messages.

Options ***policy-names***—Name of one or more import policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring PIM Bootstrap Properties
- **export (Bootstrap) on page 2091**

import (PIM)

Syntax import [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols **pim**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim**],
[edit protocols **pim**],
[edit routing-instances *routing-instance-name* protocols **pim**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply one or more policies to routes being imported into the routing table from PIM. Use the **import** statement to filter PIM join messages from entering the network.

Options ***policy-names***—Name of one or more policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Filtering Incoming PIM Join Messages

infinity

Syntax infinity [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols **pim spt-threshold**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim spt-threshold**],
[edit protocols **pim spt-threshold**],
[edit routing-instances *routing-instance-name* protocols **pim spt-threshold**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply one or more policies to set the SPT threshold to infinity for a source-group address pair. Use the **infinity** statement to prevent the last-hop routing device from transitioning from the RPT rooted at the RP to an SPT rooted at the source for that source-group address pair.

Options ***policy-names***—Name of one or more policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring the PIM SPT Threshold Policy

install

Syntax install;

Hierarchy Level [edit protocols **igmp-snooping vlan** *vlan-id* | *vlan-number* **data-forwarding receiver**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Install forwarding entries in the multicast receiver VLAN. By default, only the multicast VLAN (MVLAN) installs forwarding entries for MVLAN groups.

Default Disabled.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- [edit protocols] Configuration Statement Hierarchy on page 48
- Example: Configuring Multicast VLAN Registration on J-EX Series Switches on page 2058
- Configuring Multicast VLAN Registration (CLI Procedure) on page 2068

interface

Syntax

```
interface (all | interface-name) {
  accept-remote-source;
  disable;
  bfd-liveness-detection {
    authentication {
      algorithmalgorithm-name;
      key-chainkey-chain-name;
      loose-check;
    }
    detection-time {
      threshold milliseconds;
    }
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    version (0 | 1 | automatic);
  }
  family (inet | inet6) {
    disable;
  }
  hello-interval seconds;
  mode (dense | sparse | sparse-dense);
  neighbor-policy [ policy-names ];
  override-interval milliseconds;
  priority number;
  propagation-delay milliseconds;
  reset-tracking-bit;
  version version;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols **pim**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim**],
[edit protocols **pim**],
[edit routing-instances *routing-instance-name* protocols **pim**]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Enable PIM on an interface and configure interface-specific properties.

Options

interface-name—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify **all**. For details about specifying interfaces, see the *Junos OS Network Interfaces Configuration Guide*.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Junos OS Multicast Protocols Configuration Guide*

interface

Syntax

```
interface interface-name {
  group-limit limit;
  multicast-router-interface;
  static {
    group ip-address;
  }
}
```

Hierarchy Level [edit protocols igmp-snooping vlan *vlan-id* | *vlan-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable IGMP snooping on an interface and configure interface-specific properties.

The remaining statements are explained separately.

Default None.

Options ***interface-name***—Name of the interface.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **show igmp-snooping vlans on page 2187**
- Example: Configuring IGMP Snooping on J-EX Series Switches on page 2055
- Configuring IGMP Snooping (CLI Procedure) on page 2063

interface

Syntax

```
interface interface-name {
  disable;
  (accounting | no-accounting);
  group-policy [ policy-names ];
  immediate-leave;
  oif-map map-name;
  passive;
  promiscuous-mode;
  ssm-map ssm-map-name;
  static {
    group multicast-group-address {
      exclude;
      group-count number;
      group-increment increment;
      source ip-address {
        source-count number;
        source-increment increment;
      }
    }
  }
  version version;
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp],
[edit protocols igmp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable IGMP on an interface and configure interface-specific properties.

Options ***interface-name***—Name of the interface. Specify the full interface name, including the physical and logical address components. To configure all interfaces, you can specify **all**. For details about specifying interfaces, see the *Junos OS Network Interfaces Configuration Guide*.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Enabling IGMP

join-load-balance

Syntax join-load-balance;

Hierarchy Level [edit logical-systems *logical-system-name* protocols **pim**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim**],
[edit protocols **pim**],
[edit routing-instances *routing-instance-name* protocols **pim**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable load balancing of PIM join messages across interfaces and routing devices.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM Join Load Balancing
- clear pim join-distribution in the *Protocols Command Reference*

local

Syntax

```
local {
  disable;
  address address;
  family (inet | inet6) {
    disable;
    address address;
    anycast-pim {
      local-address address;
        rp-set {
        address address <forward-msdp-sa>;
      }
    }
    group-ranges {
      destination-ip-prefix</prefix-length>;
    }
    hold-time seconds;
    priority number;
  }
  group-ranges {
    destination-ip-prefix</prefix-length>;
  }
  hold-time seconds;
  priority number;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols **pim rp**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim rp**],
[edit protocols **pim rp**],
[edit routing-instances *routing-instance-name* protocols **pim rp**]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.
The remaining statements are explained separately.

Description

Configure the routing device's RP properties.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Local PIM RPs

local-address

Syntax local-address *address*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols **pim rp local family** (inet | inet6) **anycast-pim**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim rp local family** (inet | inet6) **anycast-pim**],
[edit protocols **pim rp local family** (inet | inet6) **anycast-pim**],
[edit routing-instances *routing-instance-name* protocols **pim rp local family** (inet | inet6) **anycast-pim**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the routing device's local address for anycast rendezvous point (RP). If this statement is omitted, the router ID is used as this address.

Options ***address***—Anycast RP IPv4 or IPv6 address, depending on **family** configuration.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring PIM Anycast with MSDP

mapping-agent-election

Syntax (mapping-agent-election | no-mapping-agent-election);

Hierarchy Level [edit logical-systems *logical-system-name* protocols **pim rp auto-rp**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim rp auto-rp**],
[edit protocols **pim rp auto-rp**],
[edit routing-instances *routing-instance-name* protocols **pim rp auto-rp**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the routing device's mapping announcements as a mapping agent.

Options **mapping-agent-election**—Mapping agents do not announce mappings when receiving mapping messages from a higher-addressed mapping agent.

no-mapping-agent-election—Mapping agents always announce mappings and do not perform mapping agent election.

Default: mapping-agent-election

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring PIM Auto-RP

maximum-rps

Syntax maximum-rps *limit*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols **pim rp embedded-rp**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim rp embedded-rp**],
[edit protocols **pim rp embedded-rp**],
[edit routing-instances *routing-instance-name* protocols **pim rp embedded-rp**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Limit the number of RPs that the routing device acknowledges.

Options *limit*—Number of RPs.

Range: 1 through 500

Default: 100

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring PIM Embedded RP for IPv6

mode

Syntax mode (dense | sparse | sparse-dense);

Hierarchy Level [edit logical-systems *logical-system-name* protocols **pim interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim interface** *interface-name*],
[edit protocols **pim interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols **pim interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure PIM to operate in sparse, dense, or sparse-dense mode.

Options

dense—Operate in dense mode.

sparse—Operate in sparse mode.

sparse-dense—Operate in sparse-dense mode.

Default: sparse

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM Dense Mode Properties
- Configuring PIM Sparse-Dense Mode Properties
- *Junos OS Multicast Protocols Configuration Guide*

multicast-router-interface

Syntax multicast-router-interface;

Hierarchy Level [edit protocols **igmp-snooping vlan** *vlan-id* | *vlan-name* **interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Statically configure an interface as a switching interface toward a multicast router (the interface to receive multicast traffic).

Default Disabled.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring IGMP Snooping on J-EX Series Switches on page 2055
- Configuring IGMP Snooping (CLI Procedure) on page 2063

neighbor-policy

Syntax neighbor-policy [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols **pim interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim interface** *interface-name*],
[edit protocols **pim interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols **pim interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply a PIM interface-level policy to filter neighbor IP addresses.

Options ***policy-name***—Name of the policy that filters neighbor IP addresses. For details about configuring policy statements, see the *Junos OS Policy Framework Configuration Guide*.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Interface-Level PIM Neighbor Policies

pim

Syntax

```
pim {
    disable;
    assert-timeout seconds;
    dense-groups {
        addresses;
    }
    dr-election-on-p2p;
    export;
    family (inet | inet6) {
        disable;
    }
    graceful-restart {
        disable;
        restart-duration seconds;
    }
    import [ policy-names ];
    interface interface-name {
        accept-remote-source;
        disable;
        bfd-liveness-detection {
            authentication {
                algorithm algorithm-name;
                key-chain key-chain-name;
                loose-check;
            }
            detection-time {
                threshold milliseconds;
            }
            minimum-interval milliseconds;
            minimum-receive-interval milliseconds;
            multiplier number;
            version (0 | 1 | automatic);
        }
        family (inet | inet6) {
            disable;
        }
        hello-interval seconds;
        mode (dense | sparse | sparse-dense);
        neighbor-policy [ policy-names ];
        override-interval milliseconds;
        priority number;
        propagation-delay milliseconds;
        reset-tracking-bit;
        version version;
    }
    join-load-balance;
    join-prune-timeout;
    nonstop-routing;
    override-interval milliseconds;
    propagation-delay milliseconds;
    reset-tracking-bit;
    rib-group group-name;
```

```
rp {
  auto-rp {
    (announce | discovery | mapping);
    (mapping-agent-election | no-mapping-agent-election);
  }
  bootstrap {
    family (inet | inet6) {
      export [ policy-names ];
      import [ policy-names ];
      priority number;
    }
  }
  bootstrap-import [ policy-names ];
  bootstrap-export [ policy-names ];
  bootstrap-priority number;
  dr-register-policy [ policy-names ];
  embedded-rp {
    group-ranges {
      destination-ip-prefix</prefix-length>;
    }
    maximum-rps limit;
  }
  local {
    family (inet | inet6) {
      address address;
      anycast-pim {
        rp-set {
          address address <forward-msdp-sa>;
        }
        disable;
        local-address address;
      }
      group-ranges {
        destination-ip-prefix</prefix-length>;
      }
      hold-time seconds;
      priority number;
    }
  }
  rp-register-policy [ policy-names ];
  spt-threshold {
    infinity [ policy-names ];
  }
  static {
    address address {
      version version;
      group-ranges {
        destination-ip-prefix</prefix-length>;
      }
    }
  }
}
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
```

```
    tunnel-devices [ mt-fpc/pic/port ];
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols],
[edit protocols],
[edit routing-instances *routing-instance-name* protocols]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Enable PIM on the routing device.

The statements are explained separately.

Default

PIM is disabled on the routing device.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM Dense Mode Properties
- Configuring PIM Dense Mode Properties
- *Junos OS Multicast Protocols Configuration Guide*

priority (Bootstrap)

Syntax

priority *number*;

Hierarchy Level

[edit logical-systems *logical-system-name* protocols **pim rp bootstrap** (inet | inet6)],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim rp bootstrap** (inet | inet6)],
[edit protocols **pim rp bootstrap** (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols **pim rp bootstrap** (inet | inet6)]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure the routing device's likelihood to be elected as the bootstrap router.

Options

number—Routing device's priority for becoming the bootstrap router. A higher value corresponds to a higher priority.

Range: 0 through a 32-bit number

Default: 0 (The routing device has the least likelihood of becoming the bootstrap router and sends packets with a priority of 0.)

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM Bootstrap Properties
- **bootstrap-priority on page 2086**

priority (PIM Interfaces)

Syntax priority *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols **pim interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim interface** *interface-name*],
[edit protocols **pim interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols **pim interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the routing device's likelihood to be elected as the designated router.

Options ***number***—Routing device's priority for becoming the designated router. A higher value corresponds to a higher priority.

Range: 1 through a 32-bit number

Default: 1 (The routing device has the least likelihood of becoming the designated router.)

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Interface Priority to Become the PIM Designated Router

priority (PIM RPs)

Syntax priority *number*;

Hierarchy Level [edit protocols **pim rp local family** (inet | inet6)],
[edit routing-instances *routing-instance-name* protocols **pim rp local family** (inet | inet6)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure this routing device's priority for becoming an RP. The bootstrap router uses this field when selecting the list of candidate RPs to send in the bootstrap message. A smaller number increases the likelihood that the routing device becomes the RP for local multicast groups. A priority value of 0 means that bootstrap router can override the group range being advertised by the candidate RP.

Options ***number***—Routing device's priority for becoming an RP. A lower value corresponds to a higher priority.

Range: 0 through 255

Default: 1

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Local PIM RPs

promiscuous-mode

Syntax promiscuous-mode;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp **interface** *interface-name*],
[edit protocols igmp **interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify that the interface accepts IGMP reports from hosts on any subnetwork.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Accepting IGMP Messages from Remote Subnetworks

proxy

Syntax proxy source-address *source-address*;

Hierarchy Level [edit protocols **igmp-snooping vlan** *vlan-id* | *vlan-number*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify that the VLAN operates in proxy mode. The proxy option is only accepted for a VLAN acting as a data-forwarding source.

Default Disabled.

Options **source-address *source-address***—IP address of the source VLAN to act as proxy.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [edit protocols] Configuration Statement Hierarchy on page 48
- Example: Configuring Multicast VLAN Registration on J-EX Series Switches on page 2058
- Configuring Multicast VLAN Registration (CLI Procedure) on page 2068

query-interval

Syntax query-interval *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp],
[edit protocols igmp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify how often the querier router sends general host-query messages.

Options ***seconds***—Time interval.
Range: 1 through 1024
Default: 125 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Modifying the IGMP Host-Query Message Interval
- **query-last-member-interval on page 2120**
- **query-response-interval on page 2120**

query-last-member-interval

Syntax query-last-member-interval *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp],
[edit protocols igmp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify how often the querier router sends group-specific query messages.

Options *seconds*—Time interval, in fractions of a second or seconds.

Range: 0.1 through 0.9, then in 1-second intervals 1 through 1024

Default: 1 second

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Modifying the IGMP Last-Member Query Interval
- **query-interval on page 2119**
- **query-response-interval on page 2120**

query-response-interval

Syntax query-response-interval *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp],
[edit protocols igmp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify how long the querier router waits to receive a response to a host-query message from a host.

Options *seconds*—The query response interval must be less than the query interval.

Range: 1 through 1024

Default: 10 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Modifying the IGMP Query Response Interval
- **query-interval on page 2119**
- **query-last-member-interval on page 2120**

receiver

Syntax

```
receiver {
  source-vlans vlan-list;
  install;
}
```

Hierarchy Level [edit protocols **igmp-snooping vlan** *vlan-id* | *vlan-number* **data-forwarding**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a VLAN as a multicast receiver VLAN of the multicast VLAN (MVLAN).

The remaining statements are explained separately.

Default Disabled.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [edit protocols] Configuration Statement Hierarchy on page 48
- Example: Configuring Multicast VLAN Registration on J-EX Series Switches on page 2058
- Configuring Multicast VLAN Registration (CLI Procedure) on page 2068

restart-duration

Syntax restart-duration *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols **pim graceful-restart**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim graceful-restart**],
[edit protocols **pim graceful-restart**],
[edit routing-instances *routing-instance-name* protocols **pim graceful-restart**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the duration of the graceful restart interval.

Options ***seconds***—Time the routing device waits (in seconds) to complete PIM sparse mode graceful restart.

Range: 30 through 300

Default: 60

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM Sparse Mode Graceful Restart

rib-group

Syntax rib-group *group-name*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols **pim**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim**],
[edit protocols **pim**],
[edit routing-instances *routing-instance-name* protocols **pim**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Associate a routing table group with PIM.

Options *group-name*—Name of the routing table group. The name must be one that you defined with the **rib-group** statement at the **[edit routing-options]** hierarchy level.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring a PIM RPF Routing Table

robust-count

Syntax robust-count *number*;

Hierarchy Level [edit protocols **igmp-snooping vlan** *vlan-id* | *vlan-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the number of intervals the switch waits before removing a multicast group from the multicast forwarding table. The length of each interval is configured using the **query-interval** statement.

Default 2

Options *number* —Number of intervals the switch waits before timing out a multicast group.
Range: 2 through 10

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring IGMP Snooping on J-EX Series Switches on page 2055
- Configuring IGMP Snooping (CLI Procedure) on page 2063

robust-count

Syntax robust-count *number*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp],
[edit protocols igmp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Tune the expected packet loss on a subnet. This factor is used to calculate the group member interval, other querier present interval, and last-member query count.

Options *number*—Robustness variable.

Range: 2 through 10

Default: 2

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Modifying the IGMP Robustness Variable

rp

Syntax

```
rp {
    auto-rp {
        (announce | discovery | mapping);
        (mapping-agent-election | no-mapping-agent-election);
    }
    bootstrap {
        family (inet | inet6) {
            export [ policy-names ];
            import [ policy-names ];
            priority number;
        }
    }
    bootstrap-export [ policy-names ];
    bootstrap-import [ policy-names ];
    bootstrap-priority number;
    dr-register-policy [ policy-names ];
    embedded-rp {
        group-ranges {
            destination-ip-prefix</prefix-length>;
        }
        maximum-rps limit;
    }
    local {
        family (inet | inet6) {
            disable;
            address address;
            anycast-pim {
                rp-set {
                    address address <forward-msdp-sa>;
                }
                local-address address;
            }
            group-ranges {
                destination-ip-prefix</prefix-length>;
            }
            hold-time seconds;
            priority number;
        }
    }
    rp-register-policy [ policy-names ];
    static {
        address address {
            version version;
            group-ranges {
                destination-ip-prefix</prefix-length>;
            }
        }
    }
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols pim],

[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim**],
[edit protocols **pim**],
[edit routing-instances *routing-instance-name* protocols **pim**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the routing device as an actual or potential RP. A routing device can be an RP for more than one group.

The remaining statements are explained separately.

Default If you do not include the **rp** statement, the routing device can never become the RP.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Junos OS Multicast Protocols Configuration Guide*

rp-register-policy

Syntax rp-register-policy [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols **pim rp**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim rp**],
[edit protocols **pim rp**],
[edit routing-instances *routing-instance-name* protocols **pim rp**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply one or more policies to control incoming PIM register messages.

Options ***policy-names***—Name of one or more import policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Register Message Filtering on a PIM RP or DR
- **dr-register-policy on page 2090**

rp-set

Syntax

```
rp-set {
    address address <forward-msdp-sa>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols **pim local family** (inet | inet6) **anycast-pim**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim local family** (inet | inet6) **anycast-pim**],
[edit protocols **pim local family** (inet | inet6) **anycast-pim**],
[edit routing-instances *routing-instance-name* protocols **pim local family** (inet | inet6) **anycast-pim**]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure a set of rendezvous point (RP) addresses for anycast RP. You can configure up to 15 RPs.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring PIM Anycast with MSDP

source

Syntax

```
source {
    groups group-prefix;
}
```

Hierarchy Level

[edit protocols **igmp-snooping vlan** *vlan-number* **data-forwarding**]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure a VLAN to be a multicast source VLAN (MVLAN).

The remaining statement is explained separately.

Default

Disabled.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [edit protocols] Configuration Statement Hierarchy on page 48
- Example: Configuring Multicast VLAN Registration on J-EX Series Switches on page 2058
- Configuring Multicast VLAN Registration (CLI Procedure) on page 2068

source

Syntax

```
source ip-address {
    source-count number;
    source-increment increment;
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp **interface** *interface-name* **static group** *multicast-group-address*],
[edit protocols igmp **interface** *interface-name* **static group** *multicast-group-address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the IP version 4 (IPv4) unicast source address for the multicast group being statically configured on an interface.

Options ***ip-address***—IPv4 unicast address.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Enabling IGMP Static Group Membership

source-vlans

Syntax source-vlans *vlan-list*;

Hierarchy Level [edit protocols **igmp-snooping vlan** *vlan-id* | *vlan-number* **data-forwarding receiver**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify a list of multicast VLANs (MVLANs) from which this multicast receiver VLAN receives multicast traffic. Either all of these MVLANs must be in proxy mode or none of them can be in proxy mode.

Default Disabled.

Options ***vlan-list***—Names of the MVLANs.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [edit protocols] Configuration Statement Hierarchy on page 48
- Example: Configuring Multicast VLAN Registration on J-EX Series Switches on page 2058
- Configuring Multicast VLAN Registration (CLI Procedure) on page 2068

spt-threshold

Syntax

```
spt-threshold {
  infinity [ policy-names ];
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols **pim**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim**],
[edit protocols **pim**],
[edit routing-instances *routing-instance-name* protocols **pim**]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Set the SPT threshold to infinity for a source-group address pair. Last-hop multicast routing devices running PIM sparse mode can forward the same stream of multicast packets onto the same LAN through an RPT rooted at the RP or an SPT rooted at the source. By default, last-hop routing devices transition to a direct SPT to the source. You can configure this routing device to set the SPT transition value to infinity to prevent this transition for any source-group address pair.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring the PIM SPT Threshold Policy

ssm-map

Syntax

ssm-map *ssm-map-name*;

Hierarchy Level

[edit logical-systems *logical-system-name* protocols igmp **interface** *interface-name*],
[edit protocols igmp **interface** *interface-name*]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Apply an SSM map to an IGMP interface.

Options

ssm-map-name—Name of SSM map.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring SSM Mapping

static

Syntax

```
static {
  address address {
    group-ranges {
      destination-ip-prefix</prefix-length>;
    }
    version version;
  }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols **pim rp**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim rp**],
[edit protocols **pim rp**],
[edit routing-instances *routing-instance-name* protocols **pim rp**]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure static RP addresses. The default static RP address is **224.0.0.0/4**. To configure other addresses, include one or more **address** statements. You can configure a static RP in a logical system only if the logical system is not directly connected to a source.

For each static RP address, you can optionally specify the PIM version and the groups for which this address can be the RP. The default PIM version is version 1.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Static PIM RPs

static (IGMP Snooping)

Syntax

```
static {
  group ip-address;
}
```

Hierarchy Level [edit protocols **igmp-snooping vlan** *vlan-id* | *vlan-name* **interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Statically define multicast groups on an interface.

The remaining statement is explained separately.

Default No multicast groups are statically defined.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring IGMP Snooping on J-EX Series Switches on page 2055
- Configuring IGMP Snooping on J-EX Series Switches (CLI Procedure) on page 2063

static

Syntax

```
static {
  group multicast-group-address {
    exclude;
    group-count number;
    group-increment increment;
    source ip-address {
      source-count number;
      source-increment increment;
    }
  }
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp **interface** *interface-name*], [edit protocols igmp **interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Test multicast forwarding on an interface without a receiver host.

The remaining statements are explained separately.

Required Privilege Level routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration.

Related Documentation

- Enabling IGMP Static Group Membership

traceoptions

Syntax

```
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols **pim**],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim**],
[edit protocols **pim**],
[edit routing-instances *routing-instance-name* protocols **pim**]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure PIM tracing options.

To specify more than one tracing operation, include multiple **flag** statements.

Default

The default PIM trace options are those inherited from the routing protocol's **traceoptions** statement included at the **[edit routing-options]** hierarchy level.

Options

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place tracing output in the **pim-log** file.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you must also include the **size** statement to specify the maximum file size.

Range: 2 through 1000 files

Default: 2 files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

PIM Tracing Flags

- **assert**—Assert messages
- **bootstrap**—Bootstrap messages
- **cache**—Packets in the PIM sparse mode routing cache

- **graft**—Graft and graft acknowledgment messages
- **hello**—Hello packets
- **join**—Join messages
- **mt**—Multicast tunnel messages
- **nsr-synchronization**—Nonstop active routing (NSR) synchronization messages
- **packets**—All PIM packets
- **prune**—Prune messages
- **register**—Register and register stop messages
- **rp**—Candidate RP advertisements
- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

Default: If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Do not allow users to read the log file.

replace—(Optional) Replace an existing trace file if there is one.

Default: If you do not include this option, tracing output is appended to an existing trace file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When ***trace-file*** again reaches this size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of trace files.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 0 KB through the maximum file size supported on your system

Default: 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration.

Related Documentation

- Configuring PIM Trace Options
- Tracing DVMRP Protocol Traffic
- Tracing MSDP Protocol Traffic
- Configuring PIM Trace Options

traceoptions

Syntax

```
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable> <match
    regex>;
  flag flag (detail | disable | receive | send);
}
```

Hierarchy Level [edit protocols igmp-snooping]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define tracing operations for IGMP snooping.

Default The **traceoptions** feature is disabled by default.

Options

file *filename* —Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

files *number* —(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached (***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify gigabytes), at which point the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000

Default: 3 files

flag *flag* —Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:

- **all**—All tracing operations.
- **general**—Trace general IGMP snooping protocol events.
- **leave**—Trace leave group messages (IGMPv2 only).
- **normal**—Trace normal IGMP snooping protocol events.
- **packets**—Trace all IGMP packets.
- **policy**—Trace policy processing.
- **query**—Trace IGMP membership query messages.
- **report**—Trace membership report messages.
- **route**—Trace routing information.
- **state**—Trace IGMP state transitions.
- **task**—Trace routing protocol task processing.
- **timer**—Trace routing protocol timer processing.

match *regex* —(Optional) Refine the output to include lines that contain the regular expression.

no-world-readable—(Optional) Restricted file access to the user who created the file.

size *size* —(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **files** option.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify gigabytes

Range: 10 KB through 1 gigabytes

Default: 128 KB

world-readable—(Optional) Enable unrestricted file access.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring IGMP Snooping on J-EX Series Switches on page 2055
- Configuring IGMP Snooping (CLI Procedure) on page 2063

traceoptions

Syntax

```
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag <flag-modifier> <disable>;
}
```

Hierarchy Level [edit logical-systems *logical-system-name* protocols igmp],
[edit protocols igmp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure IGMP tracing options.

To specify more than one tracing operation, include multiple **flag** statements.

To trace the paths of multicast packets, use the **mtrace** command, as described in the *Junos OS System Basics and Services Command Reference*.

Default The default IGMP trace options are those inherited from the routing protocols **traceoptions** statement included at the **[edit routing-options]** hierarchy level.

Options **disable**—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**. We recommend that you place tracing output in the file **igmp-log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you must also include the **size** statement to specify the maximum file size.

Range: 2 through 1000 files

Default: 2 files

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

IGMP Tracing Flags

- **leave**—Leave group messages (for IGMP version 2 only).
- **mtrace**—Mtrace packets. Use the **mtrace** command to troubleshoot the software.
- **packets**—All IGMP packets.

- **query**—IGMP membership query messages, including general and group-specific queries.
- **report**—Membership report messages.

Global Tracing Flags

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Interface transactions and processing
- **timer**—Timer usage

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

Default: If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Do not allow users to read the log file.

replace—(Optional) Replace an existing trace file if there is one.

Default: If you do not include this option, tracing output is appended to an existing trace file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When ***trace-file*** again reaches this size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also include the **files** statement to specify the maximum number of trace files.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 1 MB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level
routing and trace—To view this statement in the configuration.
routing-control and trace-control—To add this statement to the configuration.

Related Documentation
- Tracing IGMP Protocol Traffic

version

Syntax
version *version*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols igmp **interface** *interface-name*],
[edit protocols igmp **interface** *interface-name*]

Release Information
Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Specify the version of IGMP.

Options
version—IGMP version number.

Range: 1, 2, or 3

Default: IGMP version 2

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Changing the IGMP Version

version (PIM)

Syntax version *version*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols **pim interface** *interface-name*],
[edit logical-systems *logical-system-name* protocols **pim rp static** address *address*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim interface** *interface-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols **pim rp static** address *address*],
[edit protocols **pim interface** *interface-name*],
[edit protocols **pim rp static** address *address*],
[edit routing-instances *routing-instance-name* protocols **pim interface** *interface-name*],
[edit routing-instances *routing-instance-name* protocols **pim rp static** address *address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the version of PIM.

Options ***version***—PIM version number.

Range: 1 or 2

Default: PIM version 2

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Changing the PIM Version

vlan

Syntax

```
vlan (vlan-id | vlan-name) {
  data-forwarding {
    source {
      groups group-prefix;
    }
    receiver {
      source-vlans vlan-list;
      install ;
    }
  }
  disable {
    interface interface-name;
  }
  immediate-leave;
  interface interface-name {
    group-limit limit;
    multicast-router-interface;
    static {
      group ip-address;
    }
  }
  proxy ;
  query-interval seconds;
  query-last-member-interval seconds;
  query-response-interval seconds;
  robust-count number;
}
```

Hierarchy Level [edit protocols igmp-snooping]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure IGMP snooping parameters for a VLAN.

The remaining statements are explained separately.

TIP: To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type ? after vlan or vlans in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range. For IGMP snooping, secondary private VLANs are not listed.

Default IGMP snooping options apply to the specified VLAN.

Options ***vlan-id***—Numeric tag for a VLAN.

Range: 0 through 4095. Tags 0 and 4095 are reserved by the Junos OS, and you should not configure them.

vlan-name—Name of a VLAN.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring IGMP Snooping (CLI Procedure) on page 2063
- IGMP Snooping on J-EX Series Switches Overview on page 2047

CHAPTER 80

Operational Mode Commands for IGMP Snooping and Multicast

clear igmp membership

Syntax

clear igmp membership
<group *address-range*>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)

clear igmp membership
<group *address-range*>
<interface *interface-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Clear Internet Group Management Protocol (IGMP) group members.

Options

none—Clear all IGMP members on all interfaces and for all address ranges.

group *address-range*—(Optional) Clear all IGMP members that are in a particular address range. An example of a range is **224.2/16**. If you omit the destination prefix length, the default is **/32**.

interface *interface-name*—(Optional) Clear all IGMP group members on an interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Related Documentation

- **show igmp group on page 2171**
- **show igmp interface on page 2175**

List of Sample Output

clear igmp membership on page 2144
clear igmp membership interface on page 2145
clear igmp membership group on page 2145

Output Fields

See **show igmp group** for an explanation of output fields.

clear igmp membership

The following sample output displays IGMP group information before and after the **clear igmp membership** command is entered:

```
user@host> show igmp group
Interface        Group           Last Reported   Timeout
so-0/0/0         224.2.127.253   10.1.128.1          186
so-0/0/0         224.2.127.254   10.1.128.1          186
so-0/0/0         239.255.255.255 10.1.128.1          187
so-0/0/0         224.1.127.255   10.1.128.1          188
local            224.0.0.6       (null)                0
local            224.0.0.5       (null)                0
local            224.2.127.254   (null)                0
local            239.255.255.255 (null)                0
local            224.0.0.2       (null)                0
```

```
local            224.0.0.13      (null)                0

user@host> clear igmp membership
Clearing Group Membership Info for so-0/0/0
Clearing Group Membership Info for so-1/0/0
Clearing Group Membership Info for so-2/0/0

user@host> show igmp group
Interface        Group           Last Reported   Timeout
local            224.0.0.6       (null)                0
local            224.0.0.5       (null)                0
local            224.2.127.254   (null)                0
local            239.255.255.255 (null)                0
local            224.0.0.2       (null)                0
local            224.0.0.13      (null)                0
```

clear igmp membership interface

The following sample output displays IGMP group information before and after the **clear igmp membership interface** command is issued:

```
user@host> show igmp group
Interface        Group           Last Reported   Timeout
so-0/0/0         224.2.127.253   10.1.128.1          210
so-0/0/0         239.255.255.255 10.1.128.1          210
so-0/0/0         224.1.127.255   10.1.128.1          215
so-0/0/0         224.2.127.254   10.1.128.1          216
local            224.0.0.6       (null)                0
local            224.0.0.5       (null)                0
local            224.2.127.254   (null)                0
local            239.255.255.255 (null)                0
local            224.0.0.2       (null)                0
local            224.0.0.13      (null)                0

user@host> clear igmp membership interface so-0/0/0
Clearing Group Membership Info for so-0/0/0

user@host> show igmp group
Interface        Group           Last Reported   Timeout
local            224.0.0.6       (null)                0
local            224.0.0.5       (null)                0
local            224.2.127.254   (null)                0
local            239.255.255.255 (null)                0
local            224.0.0.2       (null)                0
local            224.0.0.13      (null)                0
```

clear igmp membership group

The following sample output displays IGMP group information before and after the **clear igmp membership group** command is entered:

```
user@host> show igmp group
Interface        Group           Last Reported   Timeout
so-0/0/0         224.2.127.253   10.1.128.1          210
so-0/0/0         239.255.255.255 10.1.128.1          210
so-0/0/0         224.1.127.255   10.1.128.1          215
so-0/0/0         224.2.127.254   10.1.128.1          216
local            224.0.0.6       (null)                0
local            224.0.0.5       (null)                0
local            224.2.127.254   (null)                0
local            239.255.255.255 (null)                0
```

```
local            224.0.0.2       (null)                0
local            224.0.0.13      (null)                0

user@host> clear igmp membership group 239.225/16
Clearing Group Membership Range 239.225.0.0/16 on so-0/0/0
Clearing Group Membership Range 239.225.0.0/16 on so-1/0/0
Clearing Group Membership Range 239.225.0.0/16 on so-2/0/0
```

```
user@host> show igmp group
Interface        Group           Last Reported   Timeout
so-0/0/0         224.1.127.255   10.1.128.1          231
so-0/0/0         224.2.127.254   10.1.128.1          233
so-0/0/0         224.2.127.253   10.1.128.1          236
local            224.0.0.6       (null)                0
local            224.0.0.5       (null)                0
local            224.2.127.254   (null)                0
local            239.255.255.255 (null)                0
local            224.0.0.2       (null)                0
local            224.0.0.13      (null)                0
```

clear igmp statistics

Syntax

clear igmp statistics
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)

clear igmp statistics
<interface *interface-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Clear Internet Group Management Protocol (IGMP) statistics.

Options

none—Clear IGMP statistics on all interfaces.

interface *interface-name*—(Optional) Clear IGMP statistics for the specified interface only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Related Documentation

- **show igmp statistics on page 2178**

List of Sample Output

clear igmp statistics on page 2148

Output Fields

See **show igmp statistics** for an explanation of output fields.

clear igmp statistics

The following sample output displays IGMP statistics information before and after the **clear igmp statistics** command is entered:

```
user@host> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type      Received       Sent  Rx errors
Membership Query           8883        459          0
V1 Membership Report          0          0          0
DVMRP                     19784      35476          0
PIM V1                    18310          0          0
Cisco Trace                   0          0          0
V2 Membership Report          0          0          0
Group Leave                   0          0          0
Mtrace Response               0          0          0
Mtrace Request                0          0          0
Domain Wide Report            0          0          0
V3 Membership Report          0          0          0
Other Unknown types                              0
IGMP v3 unsupported type                         0
IGMP v3 source required for SSM                  0
IGMP v3 mode not applicable for SSM              0

IGMP Global Statistics
Bad Length                    0
Bad Checksum                  0
```

```
Bad Receive If               0
Rx non-local              1227

user@host> clear igmp statistics
user@host> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type      Received       Sent  Rx errors
Membership Query              0          0          0
V1 Membership Report          0          0          0
DVMRP                         0          0          0
PIM V1                        0          0          0
Cisco Trace                   0          0          0
V2 Membership Report          0          0          0
Group Leave                   0          0          0
Mtrace Response               0          0          0
Mtrace Request                0          0          0
Domain Wide Report            0          0          0
V3 Membership Report          0          0          0
Other Unknown types                                  0
IGMP v3 unsupported type                             0
IGMP v3 source required for SSM                      0
IGMP v3 mode not applicable for SSM                  0
IGMP Global Statistics
Bad Length                    0
Bad Checksum                  0
Bad Receive If                0
Rx non-local                  0
```

clear igmp-snooping membership

Syntax clear igmp-snooping membership
<vlan *vlan-id* | *vlan-name*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Clear IGMP snooping membership information.

Options vlan *vlan-id* —Numeric tag identifier of the VLAN.

vlan *vlan-name* —Name of the VLAN.

Required Privilege Level view

Related Documentation
- show igmp-snooping membership on page 2181

List of Sample Output clear igmp-snooping membership on page 2150

clear igmp-snooping membership

```
user@switch> clear igmp-snooping membership vlan employee-vlan
```

clear igmp-snooping statistics

Syntax **clear igmp-snooping statistics**

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Clear IGMP snooping statistics.

Required Privilege Level view

Related Documentation
- **show igmp-snooping statistics on page 2185**

List of Sample Output **clear igmp-snooping statistics on page 2151**

clear igmp-snooping statistics

```
user@switch> clear igmp-snooping statistics
```

clear multicast bandwidth-admission

Syntax

clear multicast bandwidth-admission
<group *group-address*>
<inet | inet6>
<instance *instance-name*>
<interface *interface-name*>
<source *source-address*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Reapply IP multicast bandwidth admissions.

Options

none—Reapply multicast bandwidth admissions for all IPv4 forwarding entries in the master routing instance.

group *group-address*—(Optional) Reapply multicast bandwidth admissions for the specified group.

inet—(Optional) Reapply multicast bandwidth admission settings for IPv4 flows.

inet6—(Optional) Reapply multicast bandwidth admission settings for IPv6 flows.

instance *instance-name*—(Optional) Reapply multicast bandwidth admission settings for the specified instance. If you do not specify an instance, the command applies to the master routing instance.

interface *interface-name*—(Optional) Examines the corresponding outbound interface in the relevant entries and acts as follows:

- If the interface is congested, and it was admitted previously, it is removed.
- If the interface was rejected previously, the **clear multicast bandwidth-admission** command enables the interface to be admitted as long as enough bandwidth exists on the interface.
- If you do not specify an interface, issuing the **clear multicast bandwidth-admission** command readmits any previously rejected interface for the relevant entries as long as enough bandwidth exists on the interface.

To manually reject previously admitted outbound interfaces, you must specify the interface.

source *source-address*—(Optional) Use with the **group** option to reapply multicast bandwidth admission settings for the specified (source, group) entry.

Required Privilege Level

clear

Related Documentation

- **show multicast interface on page 2191**

List of Sample Output

clear multicast bandwidth-admission on page 2153

Output Fields When you enter this command, you are provided feedback on the status of your request.

clear multicast bandwidth-admission

```
user@host> clear multicast bandwidth-admission
```

clear multicast scope

Syntax

clear multicast scope
<inet | inet6>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)

clear multicast scope
<inet | inet6>
<interface *interface-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Clear IP multicast scope statistics.

Options

none—(Same as **logical-system all**) Clear multicast scope statistics.

inet—(Optional) Clear multicast scope statistics for IPv4 family addresses.

inet6—(Optional) Clear multicast scope statistics for IPv6 family addresses.

interface *interface-name*—(Optional) Clear multicast scope statistics on a specific interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Related Documentation

- **show multicast scope on page 2207**

List of Sample Output

clear multicast scope on page 2154

Output Fields

When you enter this command, you are provided feedback on the status of your request.

clear multicast scope

```
user@host> clear multicast scope
```

clear multicast sessions

Syntax
clear multicast sessions
<logical-system (all | *logical-system-name*)>
<*regular-expression*>

Syntax (J-EX Series Switch)
clear multicast sessions
<*regular-expression*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Clear IP multicast sessions.

Options
none—(Same as **logical-system all**) Clear multicast sessions.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

regular-expression—(Optional) Clear only multicast sessions that contain the specified regular expression.

Required Privilege Level
clear

Related Documentation
- **show multicast sessions on page 2209**

List of Sample Output
clear multicast sessions on page 2155

Output Fields
When you enter this command, you are provided feedback on the status of your request.

clear multicast sessions
```
user@host> clear multicast sessions
```

clear multicast statistics

Syntax

clear multicast statistics
<inet | inet6>
<instance *instance-name*>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)

clear multicast statistics
<inet | inet6>
<instance *instance-name*>
<interface *interface-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Clear IP multicast statistics.

Options

none—Clear multicast statistics for all supported address families on all interfaces.

inet—(Optional) Clear multicast statistics for IPv4 family addresses.

inet6—(Optional) Clear multicast statistics for IPv6 family addresses.

instance *instance-name*—(Optional) Clear multicast statistics for the specified instance.

interface *interface-name*—(Optional) Clear multicast statistics on a specific interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

clear

Related Documentation

- show multicast statistics

List of Sample Output

clear multicast statistics on page 2156

Output Fields

When you enter this command, you are provided feedback on the status of your request.

clear multicast statistics

```
user@host> clear multicast statistics
```

clear pim join

Syntax

clear pim join
<*group-address*>
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)

clear pim join
<*group-address*>
<inet | inet6>
<instance *instance-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Clear the Protocol Independent Multicast (PIM) join and prune states.

Options

none—Clear the PIM join and prune states for all groups, family addresses, and instances.

group-address—(Optional) Clear the PIM join and prune states for a group address.

inet | inet6—(Optional) Clear the PIM join and prune states for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Clear the join and prune states for a specific PIM-enabled routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information

The **clear pim join** command cannot be used to clear the PIM join and prune state on a backup Routing Engine when nonstop active routing is enabled.

Required Privilege Level

clear

Related Documentation

- **show pim join on page 2219**

List of Sample Output

clear pim join on page 2157

Output Fields

When you enter this command, you are provided feedback on the status of your request.

clear pim join

```
user@host> clear pim join
```

clear pim register

Syntax

clear pim register
<inet | inet6>
<instance *instance-name*>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)

clear pim register
<inet | inet6>
<instance *instance-name*>
<interface *interface-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Clear Protocol Independent Multicast (PIM) register message counters.

Options

none—Clear PIM register message counters for all family addresses, instances, and interfaces.

inet | inet6—(Optional) Clear PIM register message counters for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Clear register message counters for a specific PIM-enabled routing instance.

interface *interface-name*—(Optional) Clear PIM register message counters for a specific interface.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information

The **clear pim register** command cannot be used to clear the PIM register state on a backup Routing Engine when nonstop active routing is enabled.

Required Privilege Level

clear

Related Documentation

- **show pim statistics on page 2235**

List of Sample Output

clear pim register on page 2158

Output Fields

When you enter this command, you are provided feedback on the status of your request.

clear pim register

```
user@host> clear pim register
```

clear pim statistics

Syntax

clear pim statistics
<inet | inet6>
<instance *instance-name*>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)

clear pim statistics
<inet | inet6>
<instance *instance-name*>
<interface *interface-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Clear Protocol Independent Multicast (PIM) statistics.

Options

none—Clear PIM statistics for all family addresses, instances, and interfaces.

inet | inet6—(Optional) Clear PIM statistics for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Clear statistics for a specific PIM-enabled routing instance.

interface *interface-name*—(Optional) Clear PIM statistics for a specific interface.

logical-system (all | *logical-system-name)*—(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information

The **clear pim statistics** command cannot be used to clear the PIM statistics on a backup Routing Engine when nonstop active routing is enabled.

Required Privilege Level

clear

Related Documentation

- **show pim statistics on page 2235**

List of Sample Output

clear pim statistics on page 2159

Output Fields

See **show pim statistics** for an explanation of output fields.

clear pim statistics

The following sample output displays PIM statistics before and after the **clear pim statistics** command is entered:

```
user@host> show pim statistics
PIM statistics on all interfaces:
PIM Message type       Received       Sent  Rx errors
Hello                         0          0          0
Register                      0          0          0
Register Stop                 0          0          0
Join Prune                    0          0          0
Bootstrap                     0          0          0
```

```
Assert                          0          0          0
Graft                           0          0          0
Graft Ack                       0          0          0
Candidate RP                    0          0          0
V1 Query                     2111       4222          0
V1 Register                     0          0          0
V1 Register Stop                0          0          0
V1 Join Prune               14200      13115          0
V1 RP Reachability              0          0          0
V1 Assert                       0          0          0
V1 Graft                        0          0          0
V1 Graft Ack                    0          0          0
PIM statistics summary for all interfaces:
Unknown type                              0
V1 Unknown type                           0
Unknown Version                           0
Neighbor unknown                          0
Bad Length                                0
Bad Checksum                              0
Bad Receive If                            0
Rx Intf disabled                       2007
Rx V1 Require V2                          0
Rx Register not RP                        0
RP Filtered Source                        0
Unknown Reg Stop                          0
Rx Join/Prune no state                 1040
Rx Graft/Graft Ack no state               0
...

user@host> clear pim statistics
user@host> show pim statistics
PIM statistics on all interfaces:
PIM Message type       Received       Sent  Rx errors
Hello                         0          0          0
Register                      0          0          0
Register Stop                 0          0          0
Join Prune                    0          0          0
Bootstrap                     0          0          0
Assert                        0          0          0
Graft                         0          0          0
Graft Ack                     0          0          0
Candidate RP                  0          0          0
V1 Query                      1          0          0
V1 Register                   0          0          0
...
```

mtrace

Syntax mtrace *source*
<routing-instance *routing-instance-name*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display trace information about an IP multicast path.

Options *source*—Source hostname or address.

routing-instance *routing-instance-name*—(Optional) Trace a particular routing instance.

Additional Information The **mtrace** command for multicast traffic is similar to the **traceroute** command used for unicast traffic. Unlike **traceroute**, **mtrace** traces traffic backwards, from the receiver to the source.

Required Privilege Level view

List of Sample Output **mtrace source on page 2162**

Output Fields Table 259 on page 2161 describes the output fields for the **mtrace** command. Output fields are listed in the approximate order in which they appear.

Table 259: mtrace Output Fields

Field Name	Field Description
Mtrace from	IP address of the receiver.
to	IP address of the source.
via group	IP address of the multicast group (if any).
Querying full reverse path	Indicates the full reverse path query has begun.
number-of-hops	Number of hops from the source to the named router or switch.
router-name	Name of the router or switch for this hop.
address	Address of the router or switch for this hop.
protocol	Protocol used (for example, PIM).
Round trip time	Average round-trip time, in milliseconds (ms).
total ttl of	Time-to-live (TTL) threshold.

mtrace source

```
user@host> mtrace 192.1.4.2
Mtrace from 192.1.4.2 to 192.1.1.2 via group 0.0.0.0
Querying full reverse path... * *
  0  routerA.lab.mycompany.net (192.1.1.2)
 -1  routerB.lab.mycompany.net (192.1.2.2)  PIM  thresh^ 1
 -2  routerC.lab.mycompany.net (192.1.3.2)  PIM  thresh^ 1
 -3  hostA.lab.mycompany.net (192.1.4.2)
Round trip time 2 ms; total ttl of 2 required.
```

mtrace from-source

Syntax

mtrace from-source source *source*
<brief | detail>
<extra-hops *extra-hops*>
<group *group*>
<interval *interval*>
<loop>
<max-hops *max-hops*>
<max-queries *max-queries*>
<multicast-response | unicast-response>
<no-resolve>
<no-router-alert>
<response *response*>
<routing-instance *routing-instance-name*>
<ttl *ttl*>
<wait-time *wait-time*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display trace information about an IP multicast path from a source to this router or switch. If you specify a group address with this command, the Junos OS returns additional information, such as packet rates and losses.

Options

brief | detail—(Optional) Display the specified level of output.

extra-hops *extra-hops*—(Optional) Number of hops to take after reaching a nonresponsive router. You can specify a number between **0** and **255**.

group *group*—(Optional) Group address for which to trace the path. The default group address is **0.0.0.0**.

interval *interval*—(Optional) Number of seconds to wait before gathering statistics again. The default value is **10** seconds.

loop—(Optional) Loop indefinitely, displaying rate and loss statistics.

max-hops *max-hops*—(Optional) Maximum hops to trace toward source. The range of values is **0** through **255**. The default value is **32** hops.

max-queries *max-queries*—(Optional) Maximum number of query attempts for any hop. The range of values is **1** through **32**. The default is **3**.

multicast-response—(Optional) Always request the response using multicast.

no-resolve—(Optional) Do not attempt to display addresses symbolically.

no-router-alert—(Optional) Do not use the router-alert IP option.

response *response*—(Optional) Send trace response to a host or multicast address.

routing-instance *routing-instance-name*—(Optional) Trace a particular routing instance.

source *source*—Source hostname or address.

ttl *ttl*—(Optional) IP time-to-live (TTL) value. You can specify a number between **0** and **255**. Local queries to the multicast group use a value of **1**. Otherwise, the default value is **127**.

unicast-response—(Optional) Always request the response using unicast.

wait-time *wait-time*—(Optional) Number of seconds to wait for a response. The default value is **3**.

Required Privilege Level view

List of Sample Output **mtrace from-source on page 2165**

Output Fields Table 260 on page 2164 describes the output fields for the **mtrace from-source** command. Output fields are listed in the approximate order in which they appear.

Table 260: mtrace from-source Output Fields

Field Name	Field Description
Mtrace from	IP address of the receiver.
to	IP address of the source.
via group	IP address of the multicast group (if any).
Querying full reverse path	Indicates the full reverse path query has begun.
number-of-hops	Number of hops from the source to the named router or switch.
router-name	Name of the router or switch for this hop.
address	Address of the router or switch for this hop.
protocol	Protocol used (for example, PIM).
Round trip time	Average round-trip time, in milliseconds (ms).
total ttl of	Time-to-live (TTL) threshold.
source	Source address.
Response Dest	Response destination address.
Overall	Average packet rate for all traffic at each hop.
Packet Statistics for Traffic From	Number of packets lost, number of packets sent, percentage of packets lost, and average packet rate at each hop.

Table 260: mtrace from-source Output Fields *(continued)*

Field Name	Field Description
Receiver	IP address receiving the multicast.
Query source	IP address sending the mtrace query.

mtrace from-source

```
user@host> mtrace from-source source 192.1.4.2 group 225.1.1.1
Mtrace from 192.1.4.2 to 192.1.1.2 via group 225.1.1.1
Querying full reverse path... * *
  0  routerA.lab.mycompany.net (192.1.1.2)
 -1  routerB.lab.mycompany.net (192.1.2.2)  PIM  thresh^ 1
 -2  routerC.lab.mycompany.net (192.1.3.2)  PIM  thresh^ 1
 -3  hostA.lab.mycompany.net (192.1.4.2)
Round trip time 2 ms; total ttl of 2 required.

Waiting to accumulate statistics...Results after 10 seconds:

Source        Response Dest    Overall     Packet Statistics For Traffic From
192.1.4.2 192.1.1.2    Packet      192.1.4.2 To 225.1.1.1
     v       __/  rtt    2 ms     Rate       Lost/Sent = Pct  Rate
192.1.2.1
192.1.3.2  routerC.lab.mycompany.net
     v     ^      ttl    2                     0/0    = --    0 pps
192.1.4.1
192.1.2.2   routerB.lab.mycompany.net
     v      \__   ttl    3                     ?/0            0 pps
192.1.1.2   192.1.1.2
  Receiver      Query Source
```

mtrace monitor

Syntax mtrace monitor

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Listen passively for IP multicast responses.To exit **mtrace monitor**, type Ctrl+c.

Options none—Trace the master instance.

Required Privilege Level view

List of Sample Output **mtrace monitor on page 2167**

Output Fields Table 261 on page 2166 describes the output fields for the **mtrace monitor** command. Output fields are listed in the approximate order in which they appear.

Table 261: mtrace monitor Output Fields

Field Name	Field Description
Mtrace query at	Date and time of the query.
by	Address of the host issuing the query.
resp to	Response destination.
qid	Query ID number.
packet from...to	IP address of the query source and default group destination.
from...to	IP address of the multicast source and the response address.
via group	IP address of the group to trace.
mxhop	Maximum hop setting.

mtrace monitor

```
user@host> mtrace monitor
Mtrace query at Oct 22 13:36:14 by 192.1.3.2, resp to 224.0.1.32, qid 74a5b8
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)

Mtrace query at Oct 22 13:36:17 by 192.1.3.2, resp to 224.0.1.32, qid 1d07ba
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)

Mtrace query at Oct 22 13:36:20 by 192.1.3.2, resp to same, qid 2fea1d
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)

Mtrace query at Oct 22 13:36:30 by 192.1.3.2, resp to same, qid 7c88ad
packet from 192.1.3.2 to 224.0.0.2
from 192.1.3.2 to 192.1.3.38 via group 224.1.1.1 (mxhop=60)
```

mtrace to-gateway

Syntax

mtrace to-gateway gateway *gateway*
<brief | detail>
<extra-hops *extra-hops*>
<group *group*>
<interface *interface-name*>
<interval *interval*>
<loop>
<max-hops *max-hops*>
<max-queries *max-queries*>
<multicast-response | unicast-response>
<no-resolve>
<no-router-alert>
<response *response*>
<routing-instance *routing-instance-name*>
<ttl *ttl*>
<unicast-response>
<wait-time *wait-time*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display trace information about a multicast path from this router or switch to a gateway router or switch.

Options

gateway *gateway*—Send the trace query to a gateway multicast address.

brief | detail—(Optional) Display the specified level of output.

extra-hops *extra-hops*—(Optional) Number of hops to take after reaching a nonresponsive router or switch. You can specify a number between **0** and **255**.

group *group*—(Optional) Group address for which to trace the path. The default group address is **0.0.0.0**.

interface *interface-name*—(Optional) Source address for sending the trace query.

interval *interval*—(Optional) Number of seconds to wait before gathering statistics again. The default value is **10**.

loop—(Optional) Loop indefinitely, displaying rate and loss statistics.

max-hops *max-hops*—(Optional) Maximum hops to trace toward the source. You can specify a number between **0** and **255**.. The default value is **32**.

max-queries *max-queries*—(Optional) Maximum number of query attempts for any hop. You can specify a number between **0** and **255**.. The default value is **3**.

multicast-response—(Optional) Always request the response using multicast.

no-resolve—(Optional) Do not attempt to display addresses symbolically.

no-router-alert—(Optional) Do not use the router-alert IP option.

response *response*—(Optional) Send trace response to a host or multicast address.

routing-instance *routing-instance-name*—(Optional) Trace a particular routing instance.

ttl *ttl*—(Optional) IP time-to-live value. You can specify a number between **0** and **225**. Local queries to the multicast group use TTL **1**. Otherwise, the default value is **127**.

unicast-response—(Optional) Always request the response using unicast.

wait-time *wait-time*—(Optional) Number of seconds to wait for a response. The default value is **3**.

Required Privilege Level view

List of Sample Output **mtrace to-gateway on page 2169**

Output Fields Table 262 on page 2169 describes the output fields for the **mtrace to-gateway** command. Output fields are listed in the approximate order in which they appear.

Table 262: mtrace to-gateway Output Fields

Field Name	Field Description
Mtrace from	IP address of the receiver.
to	IP address of the source.
via group	IP address of the multicast group (if any).
Querying full reverse path	Indicates the full reverse path query has begun.
number-of-hops	Number of hops from the source to the named router or switch.
router-name	Name of the router or switch for this hop.
address	Address of the router or switch for this hop.
protocol	Protocol used (for example, PIM).
Round trip time	Average round-trip time, in milliseconds (ms).
total ttl of	Time-to-live (TTL) threshold.

mtrace to-gateway

```
user@host> mtrace to-gateway gateway 192.1.3.2 group 225.1.1.1 interface 192.1.1.73 brief

Mrace from 192.1.1.73 to 192.1.1.2 via group 225.1.1.1
Querying full reverse path... * *
  0  routerA.lab.mycompany.net (192.1.1.2)
 -1  routerA.lab.mycompany.net (192.1.1.2)  PIM  thresh^ 1
 -2  routerB.lab.mycompany.net (192.1.2.2)  PIM  thresh^ 1
```

```
 -3  routerC.lab.mycompany.net (192.1.3.2)  PIM  thresh^ 1
Round trip time 2 ms; total ttl of 3 required.
```

show igmp group

Syntax
show igmp group
<brief | detail>
<*group-name*>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show igmp group
<brief | detail>
<*group-name*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display Internet Group Management Protocol (IGMP) group membership information.

Options
none—Display standard information about membership for all IGMP groups.

brief | detail—(Optional) Display the specified level of output.

group-name—(Optional) Display group membership for the specified IP address only.

logical-system (all | *logical-system-name)*—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- **clear igmp membership on page 2144**

List of Sample Output
show igmp group (Include Mode) on page 2172
show igmp group (Exclude Mode) on page 2173
show igmp group brief on page 2173
show igmp group detail on page 2173

Output Fields
Table 263 on page 2171 describes the output fields for the **show igmp group** command. Output fields are listed in the approximate order in which they appear.

Table 263: show igmp group Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface that received the IGMP membership report. A name of **local** indicates that the local routing device joined the group itself.	All levels
Group	Group address.	All levels
Group Mode	Mode the SSM group is operating in: **Include** or **Exclude**.	All levels
Source	Source address.	All levels

Table 263: show igmp group Output Fields *(continued)*

Field Name	Field Description	Level of Output
Source timeout	Time remaining until the group traffic is no longer forwarded. The timer is refreshed when a listener in include mode sends a report. A group in exclude mode or configured as a static group displays a zero timer.	**detail**
Last reported by	Address of the host that last reported membership in this group.	All levels
Timeout	Time remaining until the group membership is removed.	**brief none**
Group timeout	Time remaining until a group in exclude mode moves to include mode. The timer is refreshed when a listener in exclude mode sends a report. A group in include mode or configured as a static group displays a zero timer.	**detail**
Type	Type of group membership: • **Dynamic**—Host reported the membership. • **Static**—Membership is configured.	All levels

show igmp group (Include Mode)

```
user@host> show igmp group
Interface: t1-0/1/0.0
    Group: 232.1.1.1
        Group mode: Include
        Source: 10.0.0.2
        Last reported by: 10.9.5.2
        Timeout:      24 Type: Dynamic
    Group: 232.1.1.1
        Group mode: Include
        Source: 10.0.0.3
        Last reported by: 10.9.5.2
        Timeout:      24 Type: Dynamic
    Group: 232.1.1.1
        Group mode: Include
        Source: 10.0.0.4
        Last reported by: 10.9.5.2
        Timeout:      24 Type: Dynamic
    Group: 232.1.1.2
        Group mode: Include
        Source: 10.0.0.4
        Last reported by: 10.9.5.2
        Timeout:      24 Type: Dynamic
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
    Group: 224.0.0.2
        Source: 0.0.0.0
        Last reported by: Local
        Timeout:       0 Type: Dynamic
    Group: 224.0.0.22
        Source: 0.0.0.0
        Last reported by: Local
        Timeout:       0 Type: Dynamic
```

show igmp group (Exclude Mode)

```
user@host> show igmp group
Interface: t1-0/1/0.0
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
    Group: 224.0.0.2
        Source: 0.0.0.0
        Last reported by: Local
        Timeout:       0 Type: Dynamic
    Group: 224.0.0.22
        Source: 0.0.0.0
        Last reported by: Local
        Timeout:       0 Type: Dynamic
```

show igmp group brief

The output for the **show igmp group brief** command is identical to that for the **show igmp group** command.

show igmp group detail

```
user@host> show igmp group detail
Interface: t1-0/1/0.0
    Group: 232.1.1.1
        Group mode: Include
        Source: 10.0.0.2
        Source timeout: 12
        Last reported by: 10.9.5.2
        Group timeout:       0 Type: Dynamic
    Group: 232.1.1.1
        Group mode: Include
        Source: 10.0.0.3
        Source timeout: 12
        Last reported by: 10.9.5.2
        Group timeout:       0 Type: Dynamic
    Group: 232.1.1.1
        Group mode: Include
        Source: 10.0.0.4
        Source timeout: 12
        Last reported by: 10.9.5.2
        Group timeout:       0 Type: Dynamic
    Group: 232.1.1.2
        Group mode: Include
        Source: 10.0.0.4
        Source timeout: 12
        Last reported by: 10.9.5.2
        Group timeout:       0 Type: Dynamic
Interface: t1-0/1/1.0
Interface: ge-0/2/2.0
Interface: ge-0/2/0.0
Interface: local
    Group: 224.0.0.2
        Group mode: Exclude
        Source: 0.0.0.0
        Source timeout: 0
        Last reported by: Local
        Group timeout:       0 Type: Dynamic
    Group: 224.0.0.22
        Group mode: Exclude
        Source: 0.0.0.0
        Source timeout: 0
```

```
Last reported by: Local
Group timeout:       0 Type: Dynamic
```

show igmp interface

Syntax
show igmp interface
<brief | detail>
<*interface-name*>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show igmp interface
<brief | detail>
<*interface-name*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display information about Internet Group Management Protocol (IGMP)-enabled interfaces.

Options
none—Display standard information about all IGMP-enabled interfaces.

brief | detail—(Optional) Display the specified level of output.

interface-name—(Optional) Display information about the specified IGMP-enabled interface only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- **clear igmp membership on page 2144**

List of Sample Output
show igmp interface on page 2177
show igmp interface brief on page 2177
show igmp interface detail on page 2177

Output Fields
Table 264 on page 2175 describes the output fields for the **show igmp interface** command. Output fields are listed in the approximate order in which they appear.

Table 264: show igmp interface Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface.	All levels
State	State of the interface: **Up** or **Down**.	All levels
Querier	Address of the routing device that has been elected to send membership queries.	All levels
Timeout	How long until the IGMP querier is declared to be unreachable, in seconds.	All levels
Version	IGMP version being used on the interface: **1** , **2** , or **3**.	All levels

Table 264: show igmp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
Groups	Number of groups on the interface.	All levels
Immediate Leave	State of the immediate leave option: • **On**—Indicates that the router removes a host from the multicast group as soon as the router receives a leave group message from a host associated with the interface. • **Off**—Indicates that after receiving a leave group message, instead of removing a host from the multicast group immediately, the router sends a group query to determine if another receiver responds.	All levels
Promiscuous Mode	State of the promiscuous mode option: • **On**—Indicates that the router can accept IGMP reports from subnetworks that are not associated with its interfaces. • **Off**—Indicates that the router can accept IGMP reports only from subnetworks that are associated with its interfaces.	All levels
Passive	State of the passive mode option: • **On**—Indicates that the router can run IGMP on the interface but not send or receive control traffic such as IGMP reports, queries, and leaves. • **Off**—Indicates that the router can run IGMP on the interface and send or receive control traffic such as IGMP reports, queries, and leaves. The **passive** statement enables you to selectively activate up to two out of a possible three available query or control traffic options. When enabled, the following options appear after the **on** state declaration: • **send-general-query**—The interface sends general queries. • **send-group-query**—The interface sends group-specific and group-source-specific queries. • **allow-receive**—The interface receives control traffic	All levels
OIF map	Name of the OIF map associated to the interface.	All levels
SSM map	Name of the source-specific multicast (SSM) map (if configured) used on the interface.	All levels
Configured Parameters	Information configured by the user: • **IGMP Query Interval**—Interval (in seconds) at which this router sends membership queries when it is the querier. • **IGMP Query Response Interval**—Time (in seconds) that the router waits for a report in response to a general query. • **IGMP Last Member Query Interval**—Time (in seconds) that the router waits for a report in response to a group-specific query. • **IGMP Robustness Count**—Number of times the router retries a query.	All levels

Table 264: show igmp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
Derived Parameters	Derived information: • **IGMP Membership Timeout**—Timeout period (in seconds) for group membership. If no report is received for these groups before the timeout expires, the group membership is removed. • **IGMP Other Querier Present Timeout**—Time (in seconds) that the router waits for the IGMP querier to send a query.	All levels

show igmp interface

```
user@host> show igmp interface
Interface: at-0/3/1.0
    Querier: 10.111.30.1
    State:         Up Timeout:    None Version:  2 Groups:      4
Interface: so-1/0/0.0
    Querier: 10.111.10.1
    State:         Up Timeout:    None Version:  2 Groups:      2
Interface: so-1/0/1.0
    Querier: 10.111.20.1
    State:         Up Timeout:    None Version:  2 Groups:      4
Immediate Leave: On
Promiscuous Mode: Off

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Derived Parameters:
IGMP Membership Timeout: 260.0
IGMP Other Querier Present Timeout: 255.0
```

show igmp interface brief

The output for the **show igmp interface brief** command is identical to that for the **show igmp interface** command. For sample output, see **show igmp interface on page 2177**.

show igmp interface detail

The output for the **show igmp interface detail** command is identical to that for the **show igmp interface** command. For sample output, see **show igmp interface on page 2177**.

show igmp statistics

Syntax
show igmp statistics
<brief | detail>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show igmp statistics
<brief | detail>
<interface *interface-name*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display Internet Group Management Protocol (IGMP) statistics.

Options
none—Display IGMP statistics for all interfaces.

brief | detail—(Optional) Display the specified level of output.

interface *interface-name*—(Optional) Display IGMP statistics about the specified interface only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- **clear igmp statistics on page 2148**

List of Sample Output
show igmp statistics on page 2179
show igmp statistics interface on page 2180

Output Fields
Table 265 on page 2178 describes the output fields for the **show igmp statistics** command. Output fields are listed in the approximate order in which they appear.

Table 265: show igmp statistics Output Fields

Field Name	Field Description
IGMP packet statistics	Heading for IGMP packet statistics for all interfaces or for the specified interface name.

Table 265: show igmp statistics Output Fields *(continued)*

Field Name	Field Description
IGMP Message type	Summary of IGMP statistics: • **Membership Query**—Number of membership queries sent and received. • **V1 Membership Report**—Number of version 1 membership reports sent and received. • **DVMRP**—Number of DVMRP messages sent or received. • **PIM V1**—Number of PIM version 1 messages sent or received. • **Cisco Trace**—Number of Cisco trace messages sent or received. • **V2 Membership Report**—Number of version 2 membership reports sent or received. • **Group Leave**—Number of group leave messages sent or received. • **Mtrace Response**—Number of Mtrace response messages sent or received. • **Mtrace Request**—Number of Mtrace request messages sent or received. • **Domain Wide Report**—Number of domain-wide reports sent or received. • **V3 Membership Report**—Number of version 3 membership reports sent or received. • **Other Unknown types**—Number of unknown message types received. • **IGMP v3 unsupported type**—Number of messages received with unknown and unsupported IGMP version 3 message types. • **IGMP v3 source required for SSM**—Number of IGMP version 3 messages received that contained no source. • **IGMP v3 mode not applicable for SSM**—Number of IGMP version 3 messages received that did not contain a mode applicable for source-specific multicast (SSM).
Received	Number of messages received.
Sent	Number of messages sent.
Rx errors	Number of received packets that contained errors.
IGMP Global Statistics	Summary of IGMP statistics for all interfaces. • **Bad Length**—Number of messages received with length errors so severe that further classification could not occur. • **Bad Checksum**—Number of messages received with a bad IP checksum. No further classification was performed. • **Bad Receive If**—Number of messages received on an interface not enabled for IGMP. • **Rx non-local**—Number of messages received from senders that are not local. • **Timed out**—Number of groups that timed out as a result of not receiving an explicit leave message. • **Rejected Report**—Number of reports dropped because of the IGMP group policy. • **Total Interfaces**—Number of interfaces configured to support IGMP.

show igmp statistics

```
user@host> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type      Received       Sent  Rx errors
Membership Query           8883        459          0
V1 Membership Report          0          0          0
DVMRP                         0          0          0
PIM V1                        0          0          0
Cisco Trace                   0          0          0
V2 Membership Report          0          0          0
```

```
Group Leave                    0          0          0
Mtrace Response                0          0          0
Mtrace Request                 0          0          0
Domain Wide Report             0          0          0
V3 Membership Report           0          0          0
Other Unknown types                                  0
IGMP v3 unsupported type                             0
IGMP v3 source required for SSM                      0
IGMP v3 mode not applicable for SSM                  0

IGMP Global Statistics
Bad Length                     0
Bad Checksum                   0
Bad Receive If                 0
Rx non-local                1227
Timed out                      0
Rejected Report                0
Total Interfaces               2
```

show igmp statistics interface

```
user@host> show igmp statistics interface fe-1/0/1.0
IGMP interface packet statistics for fe-1/0/1.0
IGMP Message type      Received       Sent  Rx errors
Membership Query              0        230          0
V1 Membership Report          0          0          0
```

show igmp-snooping membership

Syntax

show igmp-snooping membership
<**brief** | **detail**>
<**interface** *interface-name*>
<**vlan** *vlan-id* | *vlan-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display IGMP snooping membership information.

Options

none—Display general parameters.

brief | **detail**—(Optional) Display the specified level of output.

interface *interface-name*—(Optional) Display IGMP snooping information for the specified interface.

vlan *vlan-id* | *vlan-name*—(Optional) Display IGMP snooping information for the specified VLAN.

Required Privilege Level

view

Related Documentation

- **show igmp-snooping route on page 2183**
- **show igmp-snooping statistics on page 2185**
- **show igmp-snooping vlans on page 2187**
- Monitoring IGMP Snooping on page 2069
- Configuring IGMP Snooping (CLI Procedure) on page 2063
- Configuring IGMP Snooping (J-Web Procedure) on page 2064

List of Sample Output

show igmp-snooping membership on page 2182
show igmp-snooping membership detail on page 2182

Output Fields

Table 266 on page 2181 lists the output fields for the **show igmp-snooping membership** command. Output fields are listed in the approximate order in which they appear.

Table 266: show igmp-snooping membership Output Fields

Field Name	Field Description	Level of Output
VLAN	Name of the VLAN.	All
Interfaces	Interfaces that are members of the listed multicast group.	All
Tag	Numerical identifier of the VLAN.	**detail**

Table 266: show igmp-snooping membership Output Fields *(continued)*

Field Name	Field Description	Level of Output
Router interfaces	List of information about multicast router interfaces: • Name of the multicast router interface. • **static** or **dynamic**—Whether the multicast router interface is static or dynamic. • **Uptime**—For static interfaces, amount of time since the interface was configured as a multicast router interface. For dynamic interfaces, amount of time since the first query was received on interface. • **timeout**—Query timeout in seconds.	**detail**
Group	IP multicast address of the multicast group. The following information is provided for the multicast group: • Name of the interface belonging to the multicast group. • **timeout**—Time (in seconds) left until the entry for the multicast group is removed. • **Last reporter**—Last host to report membership for the multicast group. • **Receiver count**—Number of interfaces that have membership in a multicast group. • **Flags**—IGMP version of the host sending a join message. • **Include source**—Source addresses from which multicast streams are allowed based on IGMPv3 reports. Shown only for IGMPv3 joins.	**detail**

show igmp-snooping membership

```
user@switch> show igmp-snooping membership
VLAN: vlan24
    224.1.1.1      *
        Interfaces: ge-0/0/0.0
    224.1.1.100    *
        Interfaces: ge-0/0/0.0
    225.1.1.100    *
        Interfaces: ge-0/0/0.0
```

show igmp-snooping membership detail

```
user@switch> show igmp-snooping membership detail
VLAN: vlan24 Tag: 24 (Index: 3)
  Router interfaces:
    ge-0/0/8.0 dynamic Uptime: 00:08:35 timeout: 254
  Group: 224.1.1.1
    ge-0/0/0.0 timeout: 223 Receiver count: 1, Flags: <V2-hosts Static>
  Group: 224.1.1.100
    ge-0/0/0.0 timeout: 170 Last reporter: 10.10.1.10 Receiver count: 1, Flags:
    <V2-hosts>
  Group: 225.1.1.100
    ge-0/0/0.0 timeout: 168 Last reporter: 10.10.1.10 Receiver count: 1, Flags:
    <V2-hosts>
```

show igmp-snooping route

Syntax

show igmp-snooping route
<**brief** | **detail**>
<**ethernet-switching** <**brief** | **detail** | **vlan** (*vlan-id* | *vlan-name*)>>
<**inet** <**brief** | **detail** | **vlan** (*vlan-id* | *vlan-name*)>>
<**vlan** *vlan-id* | *vlan-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.
Option **inet** enhanced to support IPv6 multicast groups in Junos OS Release 10.2 for J-EX Series switches.

Description

Display IGMP snooping route information.

Options

none—Display general parameters.

brief | **detail**—(Optional) Display the specified level of output.

ethernet-switching—(Optional) Display Ethernet switching information.

inet—(Optional) Display **inet** information for IPv4 and IPv6 multicast groups. For Layer 3 IPv6 multicast routes, display information about the routing table, the routing next hop, and the Layer 2 next hop.

vlan *vlan-id* | *vlan-name*—(Optional) Display route information for the specified VLAN.

Required Privilege Level

view

Related Documentation

- **show igmp-snooping statistics on page 2185**
- **show igmp-snooping vlans on page 2187**

List of Sample Output

show igmp-snooping route on page 2184
show igmp-snooping route inet detail (IPv6 Multicast Route) on page 2184
show igmp-snooping route vlan v1 on page 2184

Output Fields

Table 267 on page 2183 lists the output fields for the **show igmp-snooping route** command. Output fields are listed in the approximate order in which they appear.

Table 267: show igmp-snooping route Output Fields

Field Name	Field Description
Table	(For internal use only. Value is always 0.)
Routing Table	(For internal use only. Value is always 0.)
VLAN	Name of the VLAN on which IGMP snooping is enabled.
Group	Multicast IPv4 or IPv6 group address.

Table 267: show igmp-snooping route Output Fields *(continued)*

Field Name	Field Description
Next-hop	ID associated with the next-hop device.
Routing next-hop	ID associated with the Layer 3 next-hop device.
Interface or **Interfaces**	Name of the interface or interfaces in the VLAN associated with the multicast group.
Layer 2 next-hop	ID associated with the Layer 2 next-hop device.

show igmp-snooping route

```
user@switch> show igmp-snooping route
VLAN             Group          Next-hop
V11              224.1.1.1, *        533
        Interfaces: ge-0/0/13.0, ge-0/0/1.0
VLAN             Group          Next-hop
v12              224.1.1.3, *        534
        Interfaces: ge-0/0/13.0, ge-0/0/0.0
```

show igmp-snooping route inet detail (IPv6 Multicast Route)

```
user@switch> show igmp-snooping route inet detail
Routing table: 0
Group: ff0e::1:ff05:1a3d, 2001::ee0:81ff:ee05:1a2e
    Routing next-hop: 587
        vlan.42
    Interface: vlan.42, VLAN: v42, Layer 2 next-hop: 506
```

show igmp-snooping route vlan v1

```
user@switch> show igmp-snooping route vlan v1
Table: 0
VLAN             Group          Next-hop
v1               224.1.1.1, *       1266
        Interfaces: ge-0/0/0.0
v1               224.1.1.3, *       1266
        Interfaces: ge-0/0/0.0
v1               224.1.1.5, *       1266
        Interfaces: ge-0/0/0.0
v1               224.1.1.7, *       1266
        Interfaces: ge-0/0/0.0
v1               224.1.1.9, *       1266
        Interfaces: ge-0/0/0.0
v1               224.1.1.11, *      1266
        Interfaces: ge-0/0/0.0
```

show igmp-snooping statistics

Syntax **show igmp-snooping statistics**

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display IGMP snooping statistics.

Required Privilege Level view

Related Documentation
- **show igmp-snooping route on page 2183**
- **show igmp-snooping vlans on page 2187**

List of Sample Output **show igmp-snooping statistics on page 2185**

Output Fields Table 268 on page 2185 lists the output fields for the **show igmp-snooping statistics** command. Output fields are listed in the approximate order in which they appear.

Table 268: show igmp-snooping statistics Output Fields

Field Name	Field Description
Bad length	IGMP packet has illegal or bad length.
Bad checksum	IGMP or IP checksum is incorrect.
Invalid interface	Packet was received through an invalid interface.
Receive unknown	Unknown IGMP type.
Timed out	Number of timeouts for all multicast groups.
IGMP Type	Type of IGMP message (Query, Report, Leave, or Other).
Received	Number of IGMP packets received.
Transmitted	Number of IGMP packets transmitted.
Recv Errors	Number of general receive errors.

show igmp-snooping statistics

```
user@switch> show igmp-snooping statistics
Bad length: 0 Bad checksum: 0 Invalid interface: 0
Not local: 0 Receive unknown: 0 Timed out: 58

    IGMP Type          Received      Transmitted    Recv Errors
    Queries:              74295                0              0
    Reports:           18148423                0       16333523
```

```
Leaves:                  0              0              0
Other:                   0              0              0
```

show igmp-snooping vlans

Syntax
show igmp-snooping vlans
<**brief** | **detail**>
<**vlan** ***vlan-id*** | *vlan-name*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display IGMP snooping VLAN information.

Options
none—Display general parameters.

brief | **detail**—(Optional) Display the specified level of output.

vlan ***vlan-id*** | **vlan** ***vlan-number***—(Optional) Display VLAN information for the specified VLAN.

Required Privilege Level
view

Related Documentation
- **show igmp-snooping route on page 2183**
- **show igmp-snooping statistics on page 2185**

List of Sample Output
show igmp-snooping vlans on page 2188
show igmp-snooping vlans vlan v10 on page 2188
show igmp-snooping vlans vlan v10 detail on page 2188

Output Fields
Table 269 on page 2187 lists the output fields for the **show igmp-snooping vlans** command. Output fields are listed in the approximate order in which they appear.

Table 269: show igmp-snooping vlans Output Fields

Field Name	Field Description	Level of Output
VLAN	Name of the VLAN.	All levels
Interfaces	Number of interfaces in the VLAN.	All levels
Groups	Number of groups in the VLAN	All levels
MRouters	Number of multicast routers associated with the VLAN.	All levels
Receivers	Number of host receivers in the VLAN.	All levels
Tag	Numerical identifier of the VLAN.	Detail
vlan-interface	Internal VLAN interface identifier.	Detail
Membership timeout	Membership timeout value.	Detail

Table 269: show igmp-snooping vlans Output Fields *(continued)*

Field Name	Field Description	Level of Output
Querier timeout	Timeout value for interfaces dynamically marked as router interfaces (interfaces that receive queries). When the querier timeout is reached, the switch marks the interface as a host interface.	Detail
Interface	Name of the interface.	Detail
Reporters	Number of dynamic groups on an interface.	Detail

show igmp-snooping vlans

```
user@switch> show igmp-snooping vlans
VLAN     Interfaces Groups MRouters Receivers
default                                   0      0        0         0
v1                                       11     50        0         0
v10                                       1      0        0         0
v11                                       1      0        0         0
v180                                      3      0        1         0
v181                                      3      0        0         0
v182                                      3      0        0         0
```

show igmp-snooping vlans vlan v10

```
user@switch> show igmp-snooping vlans vlan v10
user@switch> show igmp-snooping vlans vlan v10
VLAN                                 Interfaces Groups MRouters Receivers
v10                                           1      0        0         0
```

show igmp-snooping vlans vlan v10 detail

```
user@switch> show igmp-snooping vlans vlan v10 detail
VLAN: v10, Tag: 10, vlan-interface: vlan.10
Membership timeout: 260, Querier timeout: 255
        Interface: ge-0/0/10.0, tagged, Groups: 0, Reporters: 0
```

show multicast flow-map

Syntax
show multicast flow-map
<brief | detail>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show multicast flow-map
<brief | detail>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display configuration information about IP multicast flow maps.

Options
none—Display configuration information about IP multicast flow maps on all systems.

brief | detail—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show multicast flow-map on page 2190
show multicast flow-map detail on page 2190

Output Fields
Table 270 on page 2189 describes the output fields for the **show multicast flow-map** command. Output fields are listed in the approximate order in which they appear.

Table 270: show multicast flow-map Output Fields

Field Name	Field Description	Levels of Output
Name	Name of the flow map.	All levels
Policy	Name of the policy associated with the flow map.	All levels
Cache-timeout	Cache timeout value assigned to the flow map.	All levels
Bandwidth	Bandwidth setting associated to the flow map.	All levels
Adaptive	Whether or not adaptive mode is enabled for the flow map.	none
Flow-map	Name of the flow map.	**detail**
Adaptive Bandwidth	Whether or not adaptive mode is enabled for the flow map.	**detail**
Redundant Sources	Redundant sources defined for the same destination group.	**detail**

show multicast flow-map

```
user@host> show multicast flow-map
Instance: master
Name            Policy                  Cache timeout        Bandwidth Adaptive
map2            policy2                         never          2000000 no
map1            policy1                    60 seconds          2000000 no
```

show multicast flow-map detail

```
user@host> show multicast flow-map detail
Instance: master
Flow-map: map1
    Policy:             policy1
    Cache Timeout:      600 seconds
    Bandwidth:          2000000
    Adaptive Bandwidth: yes
    Redundant Sources:  11.11.11.11
    Redundant Sources:  11.11.11.12
    Redundant Sources:  11.11.11.13
```

show multicast interface

Syntax
show multicast interface
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show multicast interface

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display bandwidth information about IP multicast interfaces.

Options
none—Display all interfaces that have multicast configured.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show multicast interface on page 2192

Output Fields
Table 271 on page 2191 describes the output fields for the **show multicast interface** command. Output fields are listed in the approximate order in which they appear.

Table 271: show multicast interface Output Fields

Field Name	Field Description
Interface	Name of the multicast interface.
Maximum bandwidth (bps)	Maximum bandwidth setting, in bits per second, for this interface.
Remaining bandwidth (bps)	Amount of bandwidth, in bits per second, remaining on the interface.
Mapped bandwidth deduction (bps)	Amount of bandwidth, in bits per second, used by any flows that are mapped to the interface. NOTE: Adding the mapped bandwidth deduction value to the local bandwidth deduction value results in the total deduction value for the interface. This field does not appear in the output when the no QoS adjustment feature is disabled.
Local bandwidth deduction (bps)	Amount of bandwidth, in bits per second, used by any mapped flows that are traversing the interface. NOTE: Adding the mapped bandwidth deduction value to the local bandwidth deduction value results in the total deduction value for the interface. This field does not appear in the output when the no QoS adjustment feature is disabled.

Table 271: show multicast interface Output Fields *(continued)*

Field Name	Field Description
Reverse OIF mapping	State of the reverse OIF mapping feature (**on** or **off**). NOTE: This field does not appear in the output when the no QoS adjustment feature is disabled.
Reverse OIF mapping no QoS adjustment	State of the no QoS adjustment feature (**on** or **off**) for interfaces that are using reverse OIF mapping. NOTE: This field does not appear in the output when the no QoS adjustment feature is disabled.
Leave timer	Amount of time a mapped interface remains active after the last mapping ends. NOTE: This field does not appear in the output when the no QoS adjustment feature is disabled.
No QoS adjustment	State (**on**) of the no QoS adjustment feature when this feature is enabled. NOTE: This field does not appear in the output when the no QoS adjustment feature is disabled.

show multicast interface

```
user@host> show multicast interface
Interface               Maximum bandwidth (bps) Remaining bandwidth (bps)
fe-0/0/3                              10000000                          0
fe-0/0/3.210                          10000000                    -2000000
fe-0/0/3.220                         100000000                   100000000
fe-0/0/3.230                          20000000                    18000000
fe-0/0/2.200                         100000000                   100000000
```

show multicast mrinfo

Syntax show multicast mrinfo
<*host*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display configuration information about IP multicast networks, including neighboring multicast router addresses.

Options none—Display configuration information about all multicast networks.

host—(Optional) Display configuration information about a particular host. Replace *host* with a hostname or IP address.

Required Privilege Level view

List of Sample Output **show multicast mrinfo on page 2194**

Output Fields Table 272 on page 2193 describes the output fields for the **show multicast mrinfo** command. Output fields are listed in the approximate order in which they appear.

Table 272: show multicast mrinfo Output Fields

Field Name	Field Description
source-address	Query address, hostname (DNS name or IP address of the source address), and multicast protocol version or the software version of another vendor.
ip-address-1—>ip-address-2	Queried router interface address and directly attached neighbor interface address, respectively.
(*name* or *ip-address*)	Name or IP address of neighbor.

[*metric/threshold/type/flags*]	Neighbor's multicast profile:
	• ***metric***—Always has a value of **1**, because **mrinfo** queries the directly connected interfaces of a device.
	• ***threshold***—Multicast threshold time-to-live (TTL). The range of values is **0** through **255**.
	• ***type***—Multicast connection type: **pim** or **tunnel**.
	• ***flags***—Flags for this route:
	  ◦ **querier**—Queried router is the designated router for the neighboring session.
	  ◦ **leaf**—Link is a leaf in the multicast network.
	  ◦ **down**—Link status indicator.

show multicast mrinfo

```
user@host> show multicast mrinfo 10.35.4.1
10.35.4.1 (10.35.4.1) [version 12.0]:
  192.168.195.166 -> 0.0.0.0 (local) [1/0/pim/querier/leaf]
  10.38.20.1 -> 0.0.0.0 (local) [1/0/pim/querier/leaf]
  10.47.1.1 -> 10.47.1.2 (10.47.1.2) [1/5/pim]
  0.0.0.0 -> 0.0.0.0 (local) [1/0/pim/down]
```

show multicast next-hops

Syntax

show multicast next-hops
<brief | detail>
<*identifier-number*>
<inet | inet6>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)

show multicast next-hops
<brief | detail>
<*identifier-number*>
<inet | inet6>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display the entries in the IP multicast next-hop table.

Options

none—Display standard information about all entries in the multicast next-hop table for all supported address families.

brief | detail—(Optional) Display the specified level of output.

identifier-number—(Optional) Show a particular next hop by ID number. The range of values is **1** through **65,535**.

inet | inet6—(Optional) Display entries for IPv4 or IPv6 family addresses, respectively.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show multicast next-hops on page 2196
show multicast next-hops brief on page 2196
show multicast next-hops detail on page 2196

Output Fields

Table 273 on page 2195 describes the output fields for the **show multicast next-hops** command. Output fields are listed in the approximate order in which they appear.

Table 273: show multicast next-hops Output Fields

Field Name	Field Description
ID	Next-hop identifier of the prefix. The identifier is returned by the routing device's Packet Forwarding Engine.
Refcnt	Number of cache entries that are using this next hop.
KRefcount	Kernel reference count for the next hop.
Downstream interface	Interface names associated with each multicast next-hop ID.

show multicast next-hops

```
user@host> show multicast next-hops
Family: INET
ID      Refcount  KRefcount Downstream interface
262142          4          2 so-1/0/0.0
262143          2          1 mt-1/1/0.49152
262148          2          1 mt-1/1/0.32769

Family: INET6
```

show multicast next-hops brief

The output for the **show multicast next-hops brief** command is identical to that for the **show multicast next-hops** command. For sample output, see **show multicast next-hops on page 2196**.

show multicast next-hops detail

The output for the **show multicast next-hops detail** command is identical to that for the **show multicast next-hops** command. For sample output, see **show multicast next-hops on page 2196**.

show multicast pim-to-igmp-proxy

Syntax
show multicast pim-to-igmp-proxy
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show multicast pim-to-igmp-proxy

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display configuration information about PIM-to-IGMP message translation, also known as PIM-to-IGMP proxy.

Options
none—Display configuration information about PIM-to-IGMP message translation.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show multicast pim-to-igmp-proxy on page 2197

Output Fields
Table 274 on page 2197 describes the output fields for the **show multicast pim-to-igmp-proxy** command. Output fields are listed in the order in which they appear.

Table 274: show multicast pim-to-igmp-proxy Output Fields

Field Name	Field Description
Proxy state	State of PIM-to-IGMP message translation, also known as PIM-to-IGMP proxy, on the configured upstream interfaces: **enabled** or **disabled**.
interface-name	Name of upstream interface (no more than two allowed) on which PIM-to-IGMP message translation is configured.

show multicast pim-to-igmp-proxy

```
user@host> show multicast pim-to-igmp-proxy
Proxy state: enabled
ge-0/1/0.1
ge-0/1/0.2
```

show multicast pim-to-mld-proxy

Syntax
show multicast pim-to-mld-proxy
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show multicast pim-to-mld-proxy

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display configuration information about PIM-to-MLD message translation, also known as PIM-to-MLD proxy.

Options
none—Display configuration information about PIM-to-MLD message translation.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show multicast pim-to-mld-proxy on page 2198

Output Fields
Table 275 on page 2198 describes the output fields for the **show multicast pim-to-mld-proxy** command. Output fields are listed in the order in which they appear.

Table 275: show multicast pim-to-mld-proxy Output Fields

Field Name	Field Description
Proxy state	State of PIM-to-MLD message translation, also known as PIM-to-MLD proxy, on the configured upstream interfaces: **enabled** or **disabled**.
interface-name	Name of upstream interface (no more than two allowed) on which PIM-to-MLD message translation is configured.

show multicast pim-to-mld-proxy

```
user@host> show multicast pim-to-mld-proxy
Proxy state: enabled
ge-0/5/0.1
ge-0/5/0.2
```

show multicast route

Syntax

show multicast route
<brief | detail | extensive>
<active | all | inactive>
<group *group*>
<inet | inet6>
<instance *instance name*>
<logical-system (all | *logical-system-name*)>
<*regular-expression*>
<source-prefix *source-prefix*>

Syntax (J-EX Series Switch)

show multicast route
<brief | detail | extensive>
<active | all | inactive>
<group *group*>
<inet | inet6>
<instance *instance name*>
<*regular-expression*>
<source-prefix *source-prefix*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display the entries in the IP multicast forwarding table. You can display similar information with the **show route table inet.1** command.

Options

none—Display standard information about all entries in the multicast forwarding table for all routing instances.

brief | detail | extensive—(Optional) Display the specified level of output.

active | all | inactive—(Optional) Display all active entries, all entries, or all inactive entries, respectively, in the multicast forwarding table.

group *group*—(Optional) Display the cache entries for a particular group.

inet | inet6—(Optional) Display multicast forwarding table entries for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display entries in the multicast forwarding table for a specific multicast instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

regular-expression—(Optional) Display information about the multicast forwarding table entries that match a UNIX-style regular expression.

source-prefix *source-prefix*—(Optional) Display the cache entries for a particular source prefix.

Required Privilege Level view

List of Sample Output
show multicast route on page 2201
show multicast route brief on page 2201
show multicast route detail on page 2201
show multicast route extensive on page 2202

Output Fields Table 276 on page 2200 describes the output fields for the **show multicast route** command. Output fields are listed in the approximate order in which they appear.

Table 276: show multicast route Output Fields

Field Name	Field Description	Level of Output
Address family	IPv4 address family (**INET**) or IPv6 address family (**INET6**).	All levels
Group	Group address.	All levels
Source	Prefix and length of the source as it is in the multicast forwarding table.	All levels
Upstream interface	Name of the interface on which the packet with this source prefix is expected to arrive.	All levels
Downstream interface list	List of interface names to which the packet with this source prefix is forwarded.	All levels
Session description	Name of the multicast session.	**detail extensive**
Statistics	Rate at which packets are being forwarded for this source and group entry (in Kbps and pps), and number of packets that have been forwarded to this prefix. If one or more of the kilobits per second packet forwarding statistic queries fails or times out, the statistics field displays **Forwarding statistics are not available**.	**detail extensive**
Next-hop ID	Next-hop identifier of the prefix. The identifier is returned by the routing device's Packet Forwarding Engine and is also displayed in the output of the **show multicast nexthops** command.	**detail extensive**
Upstream protocol	Protocol running on the interface on which the packet with this source prefix is expected to arrive.	**detail extensive**
Route state	Whether the group is **Active** or **Inactive**.	**extensive**
Forwarding state	Whether the prefix is pruned or forwarding.	**extensive**
Cache lifetime/timeout	Number of seconds until the prefix is removed from the multicast forwarding table. A value of **never** indicates a permanent forwarding entry.	**extensive**
Wrong incoming interface notifications	Number of times that the upstream interface was not available.	**extensive**

show multicast route

```
user@host> show multicast route
Family: INET

Group: 228.0.0.0
    Source: 10.255.14.144/32
    Upstream interface: local
    Downstream interface list:
        so-1/0/0.0

Group: 239.1.1.1
    Source: 10.255.14.144/32
    Upstream interface: local
    Downstream interface list:
        so-1/0/0.0

Group: 239.1.1.1
    Source: 10.255.70.15/32
    Upstream interface: so-1/0/0.0
    Downstream interface list:
        mt-1/1/0.49152

Family: INET6
```

show multicast route brief

The output for the **show multicast route brief** command is identical to that for the **show multicast route** command. For sample output, see **show multicast route on page 2201**.

show multicast route detail

```
user@host> show multicast route detail
Family: INET

Group: 228.0.0.0
    Source: 10.255.14.144/32
    Upstream interface: local
    Downstream interface list:
        so-1/0/0.0
    Session description: Unknown
    Statistics: 8 kBps, 100 pps, 45272 packets
    Next-hop ID: 262142
    Upstream protocol: PIM

Group: 239.1.1.1
    Source: 10.255.14.144/32
    Upstream interface: local
    Downstream interface list:
        so-1/0/0.0
    Session description: Administratively Scoped
    Statistics: 0 kBps, 0 pps, 13404 packets
    Next-hop ID: 262142
    Upstream protocol: PIM

Group: 239.1.1.1
    Source: 10.255.70.15/32
    Upstream interface: so-1/0/0.0
    Downstream interface list:
        mt-1/1/0.49152
    Session description: Administratively Scoped
    Statistics: 0 kBps, 0 pps, 38 packets
    Next-hop ID: 262143
    Upstream protocol: PIM

Family: INET6
```

show multicast route extensive

```
user@host> show multicast route extensive
Family: INET

Group: 228.0.0.0
    Source: 10.255.14.144/32
    Upstream interface: local
    Downstream interface list:
        so-1/0/0.0
    Session description: Unknown
    Statistics: 8 kBps, 100 pps, 46454 packets
    Next-hop ID: 262142
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: 360 seconds
    Wrong incoming interface notifications: 0

Group: 239.1.1.1
    Source: 10.255.14.144/32
    Upstream interface: local
    Downstream interface list:
        so-1/0/0.0
    Session description: Administratively Scoped
    Statistics: 0 kBps, 0 pps, 13404 packets
    Next-hop ID: 262142
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: 348 seconds
    Wrong incoming interface notifications: 0

Group: 239.1.1.1
    Source: 10.255.70.15/32
    Upstream interface: so-1/0/0.0
    Downstream interface list:
        mt-1/1/0.49152
    Session description: Administratively Scoped
    Statistics: 0 kBps, 0 pps, 40 packets
    Next-hop ID: 262143
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: 360 seconds
    Wrong incoming interface notifications: 1

Family: INET6
```

show multicast rpf

Syntax

show multicast rpf
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<*prefix*>
<summary>

Syntax (J-EX Series Switch)

show multicast rpf
<inet | inet6>
<instance *instance-name*>
<*prefix*>
<summary>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display information about multicast reverse-path-forwarding (RPF) calculations.

Options

none—Display RPF calculation information for all supported address families.

inet | inet6—(Optional) Display the RPF calculation information for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display information about multicast RPF calculations for a specific multicast instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

prefix—(Optional) Display the RPF calculation information for the specified prefix.

summary—(Optional) Display summary of all multicast RPF information.

Required Privilege Level

view

List of Sample Output

show multicast rpf on page 2204
show multicast rpf inet6 on page 2205
show multicast rpf prefix on page 2206
show multicast rpf summary on page 2206

Output Fields Table 277 on page 2204 describes the output fields for the **show multicast rpf** command. Output fields are listed in the approximate order in which they appear.

Table 277: show multicast rpf Output Fields

Field Name	Field Description
Instance	Name of the routing instance. (Displayed when multicast is configured within a routing instance.)
Source prefix	Prefix and length of the source as it exists in the multicast forwarding table.
Protocol	How the route was learned.
Interface	Upstream RPF interface.
Neighbor	Upstream RPF neighbor.

show multicast rpf

```
user@host> show multicast rpf

Multicast RPF table: inet.0, 12 entries

0.0.0.0/0
    Protocol: Static

10.255.14.132/32
    Protocol: Direct
    Interface: lo0.0

10.255.245.91/32
    Protocol: IS-IS
    Interface: so-1/1/1.0
    Neighbor: 192.168.195.21

127.0.0.1/32
Inactive172.16.0.0/12
Protocol: Static
Interface: fxp0.0
Neighbor: 192.168.14.254

192.168.0.0/16
Protocol: Static
Interface: fxp0.0
Neighbor: 192.168.14.254

192.168.14.0/24
Protocol: Direct
Interface: fxp0.0

192.168.14.132/32
Protocol: Local

192.168.195.20/30
Protocol: Direct
Interface: so-1/1/1.0
```

```
192.168.195.22/32
Protocol: Local

192.168.195.36/30
Protocol: IS-IS
Interface: so-1/1/1.0
Neighbor: 192.168.195.21
```

show multicast rpf inet6

```
user@host> show multicast rpf inet6

Multicast RPF table: inet6.0, 12 entries

::10.255.14.132/128
    Protocol: Direct
    Interface: lo0.0

::10.255.245.91/128
Protocol: IS-IS
Interface: so-1/1/1.0
Neighbor: fe80::2a0:a5ff:fe28:2e8c

::192.168.195.20/126
Protocol: Direct
Interface: so-1/1/1.0

::192.168.195.22/128
Protocol: Local

::192.168.195.36/126
Protocol: IS-IS
Interface: so-1/1/1.0
Neighbor: fe80::2a0:a5ff:fe28:2e8c

::192.168.195.76/126
Protocol: Direct
Interface: fe-2/2/0.0

::192.168.195.77/128
Protocol: Local

fe80::/64
Protocol: Direct
Interface: so-1/1/1.0

fe80::290:69ff:fe0c:993a/128
Protocol: Local

fe80::2a0:a5ff:fe12:84f/128
Protocol: Direct
Interface: lo0.0

ff02::2/128
Protocol: PIM

ff02::d/128
Protocol: PIM
```

show multicast rpf prefix

```
user@host> show multicast rpf ff02::/16

Multicast RPF table: inet6.0, 13 entries

ff02::2/128
    Protocol: PIM

ff02::d/128
    Protocol: PIM

...
```

show multicast rpf summary

```
user@host> show multicast rpf summary

Multicast RPF table: inet.0, 16 entries
Multicast RPF table: inet6.0, 12 entries
```

show multicast scope

Syntax
show multicast scope
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show multicast scope
<inet | inet6>
<instance *instance-name*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display administratively scoped IP multicast information.

Options
none—Display standard information about administratively scoped multicast information for all supported address families in all routing instances.

inet | inet6—(Optional) Display scoped multicast information for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display administratively scoped information for a specific multicast instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show multicast scope on page 2208
show multicast scope inet on page 2208
show multicast scope inet6 on page 2208

Output Fields
Table 278 on page 2207 describes the output fields for the **show multicast scope** command. Output fields are listed in the approximate order in which they appear.

Table 278: show multicast scope Output Fields

Field Name	Field Description
Scope name	Name of the multicast scope.
Group Prefix	Range of multicast groups that are scoped.
Interface	Interface that is the boundary of the administrative scope.
Resolve Rejects	Number of kernel resolve rejects.

show multicast scope

```
user@host> show multicast scope
                                                                Resolve
Scope name           Group Prefix       Interface               Rejects
232-net              232.232.0.0/16     fe-0/0/0.1                    0
local                239.255.0.0/16     fe-0/0/0.1                    0
local                ff05::/16          fe-0/0/0.1                    0
larry                ff05::1234/128     fe-0/0/0.1                    0
```

show multicast scope inet

```
user@host> show multicast scope inet
                                                                Resolve
Scope name           Group Prefix       Interface               Rejects
232-net              232.232.0.0/16     fe-0/0/0.1                    0
local                239.255.0.0/16     fe-0/0/0.1                    0
```

show multicast scope inet6

```
user@host> show multicast scope inet6
                                                                Resolve
Scope name           Group Prefix       Interface               Rejects
local                ff05::/16          fe-0/0/0.1                    0
larry                ff05::1234/128     fe-0/0/0.1                    0
```

show multicast sessions

Syntax

show multicast sessions
<brief | detail | extensive>
<logical-system (all | *logical-system-name*)>
<*regular-expression*>

Syntax (J-EX Series Switch)

show multicast sessions
<brief | detail | extensive>
<*regular-expression*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display information about announced IP multicast sessions.

Options

none—Display standard information about all multicast sessions for all routing instances.

brief | detail | extensive—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

regular-expression—(Optional) Display information about announced sessions that match a UNIX-style regular expression.

Required Privilege Level

view

List of Sample Output

show multicast sessions on page 2210
show multicast sessions regular-expression detail on page 2210

Output Fields

Table 279 on page 2209 describes the output fields for the **show multicast sessions** command. Output fields are listed in the approximate order in which they appear.

Table 279: show multicast sessions Output Fields

Field Name	Field Description
session-name	Name of the known announced multicast sessions.

show multicast sessions

```
user@host> show multicast sessions
1-Department of Biological Sciences, LSU
  ...
  Monterey Bay - DockCam
  Monterey Bay - JettyCam
  Monterey Bay - StandCam
  Monterey DockCam
  Monterey DockCam / ROV cam
  ...
  NASA TV (MPEG-1)
  ...
  UO Broadcast - NASA Videos - 25 Years of Progress
  UO Broadcast - NASA Videos - Journey through the Solar System
  UO Broadcast - NASA Videos - Life in the Universe
  UO Broadcast - NASA Videos - Nasa and the Airplane
  UO Broadcasts OPB's Oregon Story
  UO DOD News Clips
  UO Medical Management of Biological Casualties (1)
  UO Medical Management of Biological Casualties (2)
  UO Medical Management of Biological Casualties (3)
...
  376 active sessions.
```

show multicast sessions regular-expression detail

```
user@host> show multicast sessions "NASA TV" detail
SDP Version: 0  Originated by: -@128.223.83.33
 Session: NASA TV (MPEG-1)
 Description: NASA television in MPEG-1 format, provided by Private University.
Please contact the UO if you have problems with this feed.
 Email: Your Name Here <multicast@lists.private.edu>
 Phone: Your Name Here <888/555-1212>
Bandwidth: AS:1000
Start time: permanent
Stop time: none
Attribute: type:broadcast
Attribute: tool:IP/TV Content Manager 3.4.14
Attribute: live:capture:1
Attribute: x-iptv-capture:mp1s
Media: video 54302 RTP/AVP 32 31 96 97
Connection Data: 224.2.231.45 ttl 127
Attribute: quality:8
Attribute: framerate:30
Attribute: rtpmap:96 WBIH/90000
Attribute: rtpmap:97 MP4V-ES/90000
Attribute: x-iptv-svr:video 128.223.91.191 live
Attribute: fmtp:32 type=mpeg1
Media: audio 28848 RTP/AVP 14 0 96 3 5 97 98 99 100 101 102 10 11 103 104 105 106
 Connection Data: 224.2.145.37 ttl 127
Attribute: rtpmap:96 X-WAVE/8000
Attribute: rtpmap:97 L8/8000/2
Attribute: rtpmap:98 L8/8000
Attribute: rtpmap:99 L8/22050/2
Attribute: rtpmap:100 L8/22050
Attribute: rtpmap:101 L8/11025/2
Attribute: rtpmap:102 L8/11025
Attribute: rtpmap:103 L16/22050/2
Attribute: rtpmap:104 L16/22050

  1 matching sessions.
```

show multicast usage

Syntax

show multicast usage
<brief | detail>
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)

show multicast usage
<brief | detail>
<inet | inet6>
<instance *instance-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display usage information about the 10 most active Distance Vector Multicast Routing Protocol (DVMRP) or Protocol Independent Multicast (PIM) groups.

Options

none—Display multicast usage information for all supported address families for all routing instances.

brief | detail—(Optional) Display the specified level of output.

inet | inet6—(Optional) Display usage information for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display information about the most active DVMRP or PIM groups for a specific multicast instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show multicast usage on page 2212
show multicast usage brief on page 2212
show multicast usage instance on page 2212
show multicast usage detail on page 2212

Output Fields

Table 280 on page 2211 describes the output fields for the **show multicast usage** command. Output fields are listed in the approximate order in which they appear.

Table 280: show multicast usage Output Fields

Field Name	Field Description
Instance	Name of the routing instance. (Displayed when multicast is configured within a routing instance.)
Group	Group address.

Table 280: show multicast usage Output Fields *(continued)*

Field Name	Field Description
Sources	Number of sources.
Packets	Number of packets that have been forwarded to this prefix. If one or more of the packets forwarded statistic queries fails or times out, the packets field displays **unavailable**.
Bytes	Number of bytes that have been forwarded to this prefix. If one or more of the packets forwarded statistic queries fails or times out, the bytes field displays **unavailable**.
Prefix	IP address.
/len	Prefix length.
Groups	Number of multicast groups.

show multicast usage

```
user@host> show multicast usage
Group           Sources Packets             Bytes
228.0.0.0       1       52847               4439148
239.1.1.1       2       13450               1125530


Prefix          /len Groups Packets             Bytes
10.255.14.144   /32  2      66254               5561304
10.255.70.15    /32  1      43                  3374...
```

show multicast usage brief

The output for the **show multicast usage brief** command is identical to that for the **show multicast usage** command. For sample output, see **show multicast usage on page 2212**.

show multicast usage instance

```
user@host> show multicast usage instance VPN-A
Group           Sources Packets             Bytes
224.2.127.254   1       5538                509496
224.0.1.39      1       13                  624
224.0.1.40      1       13                  624

Prefix          /len Groups Packets             Bytes
192.168.195.34  /32  1      5538                509496
10.255.14.30    /32  1      13                  624
10.255.245.91   /32  1      13                  624
...
```

show multicast usage detail

```
user@host> show multicast usage detail
Group           Sources Packets             Bytes
228.0.0.0       1       53159               4465356
    Source: 10.255.14.144   /32 Packets: 53159 Bytes: 4465356
239.1.1.1       2       13450               1125530
    Source: 10.255.14.144   /32 Packets: 13407 Bytes: 1122156
    Source: 10.255.70.15    /32 Packets: 43 Bytes: 3374


Prefix          /len Groups Packets             Bytes
```

```
10.255.14.144   /32  2      66566                5587512
    Group: 228.0.0.0       Packets: 53159 Bytes: 4465356
    Group: 239.1.1.1       Packets: 13407 Bytes: 1122156
10.255.70.15    /32  1      43                   3374
    Group: 239.1.1.1       Packets: 43 Bytes: 3374
```

show pim bootstrap

Syntax

show pim bootstrap
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)

show pim bootstrap
<instance *instance-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

For sparse mode only, display information about Protocol Independent Multicast (PIM) bootstrap routers.

Options

none—Display PIM bootstrap router information for all routing instances.

instance *instance-name*—(Optional) Display information about bootstrap routers for a specific PIM-enabled routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show pim bootstrap on page 2215
show pim bootstrap instance on page 2215

Output Fields

Table 281 on page 2214 describes the output fields for the **show pim bootstrap** command. Output fields are listed in the approximate order in which they appear.

Table 281: show pim bootstrap Output Fields

Field Name	Field Description
Instance	Name of the routing instance.
BSR	Bootstrap router.
Pri	Priority of the routing device to be elected to be the bootstrap router.
Local address	Local routing device's address.
Pri	Local routing device's address priority to be elected as the bootstrap router.
State	Local routing device's election state: **Candidate**, **Elected**, or **Ineligible**.
Timeout	How long until the local routing device declares the bootstrap router to be unreachable, in seconds.

show pim bootstrap

```
user@host> show pim bootstrap
Instance: PIM.master

BSR                     Pri Local address           Pri State      Timeout
None                      0 10.255.71.46              0 InEligible       0
feco:1:1:1:1:0:aff:785c 34 feco:1:1:1:1:0:aff:7c12   0 InEligible       0
```

show pim bootstrap instance

```
user@host> show pim bootstrap instance VPN-A
Instance: PIM.VPN-A

BSR             Pri Local address   Pri State      Timeout
None              0 192.168.196.105   0 InEligible       0
```

show pim interfaces

Syntax
show pim interfaces
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show pim interfaces
<inet | inet6>
<instance *instance-name*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display information about the interfaces on which Protocol Independent Multicast (PIM) is configured.

Options
none—Display interface information for all family addresses for all routing instances.

inet | inet6—(Optional) Display interface information for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display information about interfaces for a specific PIM-enabled routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show pim interfaces on page 2217
show pim interfaces inet on page 2218
show pim interfaces inet6 on page 2218

Output Fields
Table 282 on page 2216 describes the output fields for the **show pim interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 282: show pim interfaces Output Fields

Field Name	Field Description
Instance	Name of the routing instance.
Name	Interface name.
State	State of the interface. The state also is displayed in the **show interfaces** command.

Table 282: show pim interfaces Output Fields *(continued)*

Field Name	Field Description
Mode	PIM mode running on the interface: • **Sparse**—In sparse mode, routing devices must join and leave multicast groups explicitly. Upstream routing devices do not forward multicast traffic to this routing device unless this device has sent an explicit request (using a join message) to receive multicast traffic. • **Dense**—Unlike sparse mode, where data is forwarded only to routing devices sending an explicit request, dense mode implements a flood-and-prune mechanism, similar to DVMRP (the first multicast protocol used to support the multicast backbone). • **Sparse-Dense**—Sparse-dense mode allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as **dense** is not mapped to a rendezvous point (RP). Instead, data packets destined for that group are forwarded using PIM-Dense Mode (PIM-DM) rules. A group specified as **sparse** is mapped to an RP, and data packets are forwarded using PIM-Sparse Mode (PIM-SM) rules.
IP	Version number of the address family on the interface: **4** (IPv4) or **6** (IPv6).
V	PIM version running on the interface: **1** or **2**.
State	State of PIM on the interface: • **DR**—Designated router. • **NotDR**—Not the designated router. • **P2P**—Point to point.
NbrCnt	Number of neighbors that have been seen on the interface.
JoinCnt(sg)	Number of (s,g) join messages that have been seen on the interface.
JointCnt(*g)	Number of (*,g) join messages that have been seen on the interface.
DR address	Address of the designated router.

show pim interfaces

```
user@host> show pim interfaces
Instance: PIM.master

Name               Stat Mode       IP V State NbrCnt JoinCnt(sg) JointCnt(*g) DR
 address
fe-0/0/0.0         Up   Sparse      4 2 DR         1           1            3
10.10.10.2
fe-0/0/3.0         Up   Sparse      4 2 DR         1           1            3
20.20.20.2
lo0.0              Up   Sparse      4 2 DR         0           0            0
10.255.72.54
pe-1/2/0.32769     Up   Sparse      4 2 P2P        0           0            0
t1-0/1/0.0         Up   Sparse      4 2 P2P        1           0            0
lo0.0              Up   Sparse      6 2 DR         0           0            0
fe80::2a0:a5ff:fe5e:209
```

show pim interfaces inet

```
user@host> show pim interfaces inet
Instance: PIM.master

Name               Stat Mode       IP V State NbrCnt JoinCnt(sg) JointCnt(*g) DR
 address
fe-0/0/0.0         Up   Sparse      4 2 DR         1           1            3
10.10.10.2
fe-0/0/3.0         Up   Sparse      4 2 DR         1           1            3
20.20.20.2
lo0.0              Up   Sparse      4 2 DR         0           0            0
10.255.72.54
pe-1/2/0.32769     Up   Sparse      4 2 P2P        0           0            0
t1-0/1/0.0         Up   Sparse      4 2 P2P        1           0            0
```

show pim interfaces inet6

```
user@host> show pim interfaces inet6
Instance: PIM.master

Name               Stat Mode       IP V State NbrCnt JoinCnt(sg) JointCnt(*g) DR
 address
lo0.0              Up   Sparse      6 2 DR         0           0            0
fe80::2a0:a5ff:fe5e:209
```

show pim join

Syntax

show pim join
<brief | detail | extensive>
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<*range*>

Syntax (J-EX Series Switch)

show pim join
<brief | detail | extensive>
<inet | inet6>
<instance *instance-name*>
<*range*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display information about Protocol Independent Multicast (PIM) groups.

Options

none—Display the standard information about PIM groups for all supported family addresses for all routing instances.

brief | detail | extensive—(Optional) Display the specified level of output.

inet | inet6—(Optional) Display PIM group information for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display information about groups for the specified PIM-enabled routing instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

range—(Optional) Address range of the group, specified as ***prefix/prefix-length***.

Required Privilege Level

view

Related Documentation

- **clear pim join on page 2157**

List of Sample Output

Output Fields

Table 283 on page 2220 describes the output fields for the **show pim join** command. Output fields are listed in the approximate order in which they appear.

Table 283: show pim join Output Fields

Field Name	Field Description
Instance	Name of the routing instance.
Family	Name of the address family: **inet** (IPv4) or **inet6** (IPv6).
R	Rendezvous Point Tree
S	Sparse
W	Wildcard
Group	Group address.
Source	Multicast source: • ***** (wildcard value) • ***ipv4-address*** • ***ipv6-address***
RP	Rendezvous point for the PIM group.
Flags	PIM flags: • **dense**—Dense mode entry. • **rptree**—Entry is on the rendezvous point tree. • **sparse**—Sparse mode entry. • **spt**—Entry is on the shortest-path tree for the source. • **wildcard**—Entry is on the shared tree.
Upstream interface	RPF interface toward the source address for the source-specific state **(S, G)** or toward the rendezvous point (RP) address for the non-source-specific state **(*, G)**.
Upstream neighbor	Information about the upstream neighbor: **Direct**, **Local**, **Unknown**, or a specific IP address.
Upstream state	Information about the upstream interface: • **Join to RP**—Sending a join to the rendezvous point. • **Join to Source**—Sending a join to the source. • **Local RP**—Sending neither joins nor prunes toward the RP, because this router is the rendezvous point. • **Local Source**—Sending neither joins nor prunes toward the source, because the source is locally attached to this routing device. • **Prune to RP**—Sending a prune to the rendezvous point. • **Prune to Source**—Sending a prune to the source.

Table 283: show pim join Output Fields *(continued)*

Field Name	Field Description
Downstream neighbors	Information about downstream interfaces: • **Interface**—Interface name for the downstream neighbor. NOTE: A pseudo PIM-SM interface appears for all IGMP-only interfaces. • **Interface address**—Address of the downstream neighbor. • **State**—Information about the downstream neighbor: **join** or **prune**. • **Flags**—PIM join flags: **R (RPtree),S (Sparse)**, **W (Wildcard)**, or **zero**.
Assert Timeout	Length of time between assert cycles on downstream interface. Not displayed if assert timer is null.
Timeout	Time remaining until the downstream join state is updated (in seconds). If the downstream join state is not updated before this keepalive timer reaches zero, the entry is deleted. If there is a directly connected host, **Timeout** is **Infinity**.

show pim join

```
user@host> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
    Source: *
    RP: 10.255.14.144
    Flags: sparse,rptree,wildcard
    Upstream interface: Local

Group: 239.1.1.1
    Source: 10.255.14.144
    Flags: sparse,spt
    Upstream interface: Local

Group: 239.1.1.1
    Source: 10.255.70.15
    Flags: sparse,spt
    Upstream interface: so-1/0/0.0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join instance

```
user@host> show pim join instance VPN-A
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 235.1.1.2
    Source: *
    RP: 10.10.47.100
    Flags: sparse,rptree,wildcard
    Upstream interface: Local

Group: 235.1.1.2
    Source: 192.168.195.74
    Flags: sparse,spt
    Upstream interface: at-0/3/1.0
```

```
Group: 235.1.1.2
    Source: 192.168.195.169
    Flags: sparse
    Upstream interface: so-1/0/1.0

Instance: PIM.VPN-A Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join detail

```
user@host> show pim join detail
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
    Source: *
    RP: 10.255.14.144
    Flags: sparse,rptree,wildcard
    Upstream interface: Local

Group: 239.1.1.1
    Source: 10.255.14.144
    Flags: sparse,spt
    Upstream interface: Local

Group: 239.1.1.1
    Source: 10.255.70.15
    Flags: sparse,spt
    Upstream interface: so-1/0/0.0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join extensive

```
user@host> show pim join extensive
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 239.1.1.1
    Source: *
    RP: 10.255.14.144
    Flags: sparse,rptree,wildcard
    Upstream interface: Local
    Upstream neighbor: Local
    Upstream state: Local RP
    Downstream neighbors:
        Interface: so-1/0/0.0
            10.111.10.2 State: Join Flags: SRW Timeout: 174
        Interface: mt-1/1/0.32768
            10.10.47.100 State: Join Flags: SRW  Timeout: Infinity

Group: 239.1.1.1
    Source: 10.255.14.144
    Flags: sparse,spt
    Upstream interface: Local
    Upstream neighbor: Local
    Upstream state: Local Source, Local RP
    Keepalive timeout: 344
    Downstream neighbors:
        Interface: so-1/0/0.0
            10.111.10.2 State: Join Flags: S Timeout: 174
        Interface: mt-1/1/0.32768
            10.10.47.100 State: Join Flags: S   Timeout: Infinity
```

```
Group: 239.1.1.1
    Source: 10.255.70.15
    Flags: sparse,spt
    Upstream interface: so-1/0/0.0
    Upstream neighbor: 10.111.10.2
    Upstream state: Local RP, Join to Source
    Keepalive timeout: 344
    Downstream neighbors:
        Interface: Pseudo-GMP
            fe-0/0/0.0 fe-0/0/1.0 fe-0/0/3.0
        Interface: so-1/0/0.0 (pruned)
            10.111.10.2 State: Prune Flags: SR Timeout: 174
        Interface: mt-1/1/0.32768
            10.10.47.100 State: Join Flags: S   Timeout: Infinity

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard
```

show pim join instance extensive

```
user@host> show pim join instance VPN-A extensive
Instance: PIM.VPN-A Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 235.1.1.2
    Source: *
    RP: 10.10.47.100
    Flags: sparse,rptree,wildcard
    Upstream interface: Local
    Upstream neighbor: Local
    Upstream state: Local RP
    Downstream neighbors:
        Interface: mt-1/1/0.32768
            10.10.47.101 State: Join Flags: SRW Timeout: 156

Group: 235.1.1.2
    Source: 192.168.195.74
    Flags: sparse,spt
    Upstream interface: at-0/3/1.0
    Upstream neighbor: 10.111.30.2
    Upstream state: Local RP, Join to Source
    Keepalive timeout: 156

Group: 235.1.1.2
    Source: 192.168.195.169
    Flags: sparse
    Upstream interface: so-1/0/1.0
    Upstream neighbor: 10.111.20.2
    Upstream state: Local RP, Join to Source
    Keepalive timeout: 156
```

show pim neighbors

Syntax

show pim neighbors
<brief | detail>
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)

show pim neighbors
<brief | detail>
<inet | inet6>
<instance *instance-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display information about Protocol Independent Multicast (PIM) neighbors.

Options

none—(Same as **brief**) Display standard information about PIM neighbors for all supported family addresses for all routing instances.

brief | detail—(Optional) Display the specified level of output.

inet | inet6—(Optional) Display information about PIM neighbors for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display information about neighbors for the specified PIM-enabled routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show pim neighbors on page 2226
show pim neighbors brief on page 2226
show pim neighbors instance on page 2226
show pim neighbors detail on page 2226
show pim neighbors detail (with BFD) on page 2226

Output Fields

Table 284 on page 2224 describes the output fields for the **show pim neighbors** command. Output fields are listed in the approximate order in which they appear.

Table 284: show pim neighbors Output Fields

Field Name	Field Description	Level of Output
Instance	Name of the routing instance.	All levels
Interface	Interface through which the neighbor is reachable.	All levels
Neighbor addr	Address of the neighboring PIM routing device.	All levels

Table 284: show pim neighbors Output Fields *(continued)*

Field Name	Field Description	Level of Output
IP	IP version: **4** or **6**.	All levels
V	PIM version running on the neighbor: **1** or **2**.	All levels
Mode	PIM mode of the neighbor: **Sparse**, **Dense**, **SparseDense**, or **Unknown**. When the neighbor is running PIM version 2, this mode is always **Unknown**.	All levels
Option	Can be one or more of the following: • **B**—Bidirectional Capable. • **H**—Hello Option Holdtime. • **G**—Generation Identifier. • **P**—Hello Option DR Priority. • **L**—Hello Option LAN Prune Delay.	**brief** none
Uptime	Time the neighbor has been operational since the PIM process was last initialized, in the format ***dd:hh:mm:ss*** **ago** for less than a week and ***nwnd:hh:mm:ss*** **ago** for more than a week.	All levels
Address	Address of the neighboring PIM router.	**detail**
BFD	Status and operational state of the Bidirectional Forwarding Detection (BFD) protocol on the interface: **Enabled**, **Operational state is up**, or **Disabled**.	**detail**
Hello Option Holdtime	Time for which the neighbor is available, in seconds. The range of values is **0** through **65,535.**	**detail**
Hello Default Holdtime	Default holdtime and the time remaining if the **holdtime** option is not in the received hello message.	**detail**
Hello Option DR Priority	Designated router election priority. The range of values is **0** through **255**.	**detail**
Hello Option Generation ID	9- or 10-digit number used to tag hello messages.	**detail**
Hello Option LAN Prune Delay	Time to wait before the neighbor receives prune messages, in the format **delay** ***nnn*** **ms override** ***nnnn*** **ms**.	**detail**
Join Suppression supported	Neighbor is capable of join suppression.	**detail**
Rx Join	Information about joins received from the neighbor. • **Group**—Group addresses in the join message. • **Source**—Address of the source in the join message. • **Timeout**—Time for which the join is valid.	**detail**

show pim neighbors

```
user@host> show pim neighbors
Instance: PIM.master
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority

Interface           IP V Mode        Option      Uptime Neighbor addr
so-1/0/0.0           4 2             HPLG      00:07:10 10.111.10.2
```

show pim neighbors brief

The output for the **show pim neighbors brief** command is identical to that for the **show pim neighbors** command. For sample output, see **show pim neighbors on page 2226**.

show pim neighbors instance

```
user@host> show pim neighbors instance VPN-A
Instance: PIM.VPN-A
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority

Interface           IP V Mode        Option      Uptime Neighbor addr
at-0/3/1.0           4 2             HPLG      00:07:54 10.111.30.2
mt-1/1/0.32768       4 2             HPLG      00:07:22 10.10.47.101
so-1/0/1.0           4 2             HPLG      00:07:50 10.111.20.2
```

show pim neighbors detail

```
user@host> show pim neighbors detail
Instance: PIM.master
Interface: fe-3/0/2.0
     Address: 192.168.195.37, IPv4, PIM v2, Mode: Sparse
 Hello Option Holdtime: 65535 seconds
 Hello Option DR Priority: 1
 Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
                               Join Suppression supported
 Rx Join: Group           Source          Timeout
          225.1.1.1       192.168.195.78        0
          225.1.1.1                             0
Interface: lo0.0
     Address: 10.255.245.91, IPv4, PIM v2, Mode: Sparse
 Hello Option Holdtime: 65535 seconds
 Hello Option DR Priority: 1
 Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
                               Join Suppression supported
Interface: pd-6/0/0.32768
     Address: 0.0.0.0, IPv4, PIM v2, Mode: Sparse
 Hello Option Holdtime: 65535 seconds
 Hello Option DR Priority: 0
 Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
                               Join Suppression supported
```

show pim neighbors detail (with BFD)

```
user@host> show pim neighbors detail
Instance: PIM.master
Interface: fe-1/0/0.0
    Address: 192.168.11.1,        IPv4, PIM v2, Mode: Sparse
        Hello Option Holdtime: 65535 seconds
        Hello Option DR Priority: 1
        Hello Option Generation ID: 836607909
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms

    Address: 192.168.11.2,        IPv4, PIM v2
        BFD: Enabled, Operational state is up
        Hello Default Holdtime: 105 seconds 104 remaining
```

```
        Hello Option DR Priority: 1
        Hello Option Generation ID: 1907549685
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms

Interface: fe-1/0/1.0
    Address: 192.168.12.1,        IPv4, PIM v2
        BFD: Disabled
        Hello Default Holdtime: 105 seconds 80 remaining
        Hello Option DR Priority: 1
        Hello Option Generation ID: 1971554705
        Hello Option LAN Prune Delay: delay 500 ms override 2000 ms
```

show pim rps

Syntax

show pim rps
<brief | detail | extensive>
<*group-address*>
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)

show pim rps
<brief | detail | extensive>
<*group-address*>
<inet | inet6>
<instance *instance-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display information about Protocol Independent Multicast (PIM) rendezvous points (RPs).

Options

none—Display standard information about PIM RPs for all groups and family addresses for all routing instances.

brief | detail | extensive—(Optional) Display the specified level of output.

group-address—(Optional) Display the RPs for a particular group. If you specify a group address, the output lists the routing device that is the RP for that group.

inet | inet6—(Optional) Display information for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display information about RPs for a specific PIM-enabled routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show pim rps on page 2230
show pim rps brief on page 2230
show pim rps instance on page 2230
show pim rps extensive on page 2231
show pim rps extensive (PIM Anycast RP in Use) on page 2231

Output Fields

Table 285 on page 2228 describes the output fields for the **show pim rps** command. Output fields are listed in the approximate order in which they appear.

Table 285: show pim rps Output Fields

Field Name	Field Description	Level of Output
Instance	Name of the routing instance.	All levels

Table 285: show pim rps Output Fields *(continued)*

Field Name	Field Description	Level of Output
Family	Name of the address family: **inet** (IPv4) or **inet6** (IPv6).	All levels
RP address	Address of the rendezvous point.	All levels
Type	Type of RP: • **auto-rp**—Address of the RP known through the Auto-RP protocol. • **bootstrap**—Address of the RP known through the bootstrap router protocol (BSR). • **embedded**—Address of the RP known through an embedded RP (IPv6). • **static**—Address of RP known through static configuration.	**brief** none
Holdtime	How long to keep the RP active, with time remaining, in seconds.	All levels
Timeout	How long until the local routing device determines the RP to be unreachable, in seconds.	All levels
Groups	Number of groups currently using this RP.	All levels
Group prefixes	Addresses of groups that this RP can span.	**brief** none
Learned via	Address and method by which the RP was learned.	**detail extensive**
Time Active	How long the RP has been active, in the format ***hh*:*mm*:*ss***.	**detail extensive**
Device Index	Index value of the order in which the Junos OS finds and initializes the interface.	**detail extensive**
Subunit	Logical unit number of the interface.	**detail extensive**
Interface	Either the encapsulation or the de-encapsulation logical interface, depending on whether this routing device is a designated router (DR) facing an RP router, or is the local RP, respectively.	**detail extensive**
Group Ranges	Addresses of groups that this RP spans.	**detail extensive**
Active groups using RP	Number of groups currently using this RP.	**detail extensive**
total	Total number of active groups for this RP.	**detail extensive**

Table 285: show pim rps Output Fields *(continued)*

Field Name	Field Description	Level of Output
Register State for RP	Current register state for each group: • **Group**—Multicast group address. • **Source**—Multicast source address for which the PIM register is sent or received, depending on whether this router is a designated router facing an RP router, or is the local RP, respectively: • **First Hop**—PIM-designated routing device that sent the Register message (the source address in the IP header). • **RP Address**—RP to which the Register message was sent (the destination address in the IP header). • **State**: On the designated router: ◦ **Send**—Sending Register messages. ◦ **Probe**—Sent a null register. If a Register-Stop message does not arrive in 5 seconds, the designated router resumes sending Register messages. ◦ **Suppress**—Received a Register-Stop message. The designated router is waiting for the timer to resume before changing to **Probe** state. • On the RP: ◦ **Receive**—Receiving Register messages.	**extensive**
Anycast-PIM rpset	If anycast RP is configured, the addresses of the RPs in the set.	**extensive**
Anycast-PIM local address used	If anycast RP is configured, the local address used by the RP.	**extensive**
Anycast-PIM Register State	If anycast RP is configured, the current register state for each group: • **Group**—Multicast group address. • **Source**—Multicast source address for which the PIM register is sent or received, depending on whether this routing device is a designated router facing an RP router, or is the local RP, respectively. • **Origin**—How the information was obtained: ◦ **DIRECT**—From a local attachment ◦ **MSDP**—From the Multicast Source Discovery Protocol (MSDP) ◦ **DR**—From the designated router	**extensive**

show pim rps

```
user@host> show pim rps
Instance: PIM.master
Address family INET
RP address              Type        Holdtime Timeout Groups Group prefixes
10.255.14.144           static             0    None      1 224.0.0.0/4

Address family INET6
```

show pim rps brief The output for the **show pim rps brief** command is identical to that for the **show pim rps** command. For sample output, see **show pim rps on page 2230**.

show pim rps instance

```
user@host> show pim rps instance VPN-A
```

```
Instance: PIM.VPN-A
Address family INET
RP address              Type        Holdtime Timeout Groups Group prefixes
10.10.47.100            static             0    None      1 224.0.0.0/4

Address family INET6
```

show pim rps extensive

```
user@host> show pim rps extensive
Instance: PIM.master

Family: INET
RP: 10.255.245.91
Learned via: static configuration
Time Active: 00:05:48
Holdtime: 45 with 36 remaining
Device Index: 122
Subunit: 32768
Interface: pd-6/0/0.32768
Group Ranges:
    224.0.0.0/4, 36s remaining
Active groups using RP:
    225.1.1.1

    total 1 groups active

Register State for RP:
Group           Source          FirstHop        RP Address      State    Timeout
225.1.1.1       192.168.195.78  10.255.14.132   10.255.245.91   Receive        0
```

show pim rps extensive (PIM Anycast RP in Use)

```
user@host> show pim rps extensive
Instance: PIM.master

Family: INET
RP: 10.10.10.2
Learned via: static configuration
Time Active: 00:54:52
Holdtime: 0
Device Index: 130
Subunit: 32769
Interface: pimd.32769
Group Ranges:
    224.0.0.0/4
Active groups using RP:
    224.10.10.10

    total 1 groups active

Anycast-PIM rpset:
          10.100.111.34
          10.100.111.17
          10.100.111.55

Anycast-PIM local address used: 10.100.111.1
Anycast-PIM Register State:
Group                   Source                          Origin
224.1.1.1               10.10.95.2                      DIRECT
224.1.1.2               10.10.95.2                      DIRECT
224.10.10.10            10.10.70.1                      MSDP
224.10.10.11            10.10.70.1                      MSDP
224.20.20.1             10.10.71.1                      DR
```

```
Address family INET6

Anycast-PIM rpset:
          ab::1
          ab::2
Anycast-PIM local address used: cd::1

Anycast-PIM Register State:
Group                   Source                                  Origin
::224.1.1.1             ::10.10.95.2                            DIRECT
::224.1.1.2             ::10.10.95.2                            DIRECT
::224.20.20.1           ::10.10.71.1                            DR
```

show pim source

Syntax

show pim source
<brief | detail>
<inet | inet6>
<instance *instance-name*>
<logical-system (all | *logical-system-name*)>
<*source-prefix*>

Syntax (J-EX Series Switch)

show pim source
<brief | detail>
<inet | inet6>
<instance *instance-name*>
<*source-prefix*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display information about the Protocol Independent Multicast (PIM) source reverse path forwarding (RPF) state.

Options

none—Display standard information about the PIM RPF state for all supported family addresses for all routing instances.

brief | detail—(Optional) Display the specified level of output.

inet | inet6—(Optional) Display information for IPv4 or IPv6 family addresses, respectively.

instance *instance-name*—(Optional) Display information about the RPF state for a specific PIM-enabled routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

source-prefix—(Optional) Display the state for source RPF states in the given range.

Required Privilege Level

view

List of Sample Output

show pim source on page 2234
show pim source brief on page 2234
show pim source detail on page 2234

Output Fields

Table 286 on page 2233 describes the output fields for the **show pim source** command. Output fields are listed in the approximate order in which they appear.

Table 286: show pim source Output Fields

Field Name	Field Description
Instance	Name of the routing instance.
RPF Address	Address of the source or reverse path.

Table 286: show pim source Output Fields *(continued)*

Field Name	Field Description
Prefix/length	Prefix and prefix length for the route used to reach the RPF address.
Upstream interface	RPF interface toward the source address.
Neighbor address	Address of the RPF neighbor used to reach the source address.

show pim source

```
user@host> show pim source
Instance: PIM.master Family: INET

Source 10.255.14.144
    Prefix 10.255.14.144/32
    Upstream interface Local
    Upstream neighbor Local

Source 10.255.70.15
    Prefix 10.255.70.15/32
    Upstream interface so-1/0/0.0
    Upstream neighbor 10.111.10.2

Instance: PIM.master Family: INET6
```

show pim source brief

The output for the **show pim source brief** command is identical to that for the **show pim source** command. For sample output, see **show pim source on page 2234**.

show pim source detail

```
user@host> show pim source detail
Instance: PIM.master Family: INET

Source 10.255.14.144
    Prefix 10.255.14.144/32
    Upstream interface Local
    Upstream neighbor Local
    Active groups:228.0.0.0
        239.1.1.1
        239.1.1.1

Source 10.255.70.15
    Prefix 10.255.70.15/32
    Upstream interface so-1/0/0.0
    Upstream neighbor 10.111.10.2
    Active groups:239.1.1.1

Instance: PIM.master Family: INET6
```

show pim statistics

Syntax

show pim statistics
<inet |inet6)>
<instance *instance-name*>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)

show pim statistics
<inet |inet6)>
<instance *instance-name*>
<interface *interface-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display Protocol Independent Multicast (PIM) statistics.

Options

none—Display PIM statistics.

inet | inet6—(Optional) Display IPv4 or IPv6 PIM statistics.

instance *instance-name*—(Optional) Display statistics for a specific routing instance enabled by Protocol Independent Multicast (PIM).

interface *interface-name*—(Optional) Display statistics about the specified interface.

logical-system (all | *logical-system-name)*—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

Related Documentation

- **clear pim statistics on page 2159**

List of Sample Output

show pim statistics on page 2240

Output Fields

Table 287 on page 2235 describes the output fields for the **show pim statistics** command. Output fields are listed in the approximate order in which they appear.

Table 287: show pim statistics Output Fields

Field Name	Field Description
PIM statistics	PIM statistics for all interfaces or for the specified interface.
PIM message type	Message type for which statistics are displayed.
Received	Number of received statistics.
Sent	Number of messages sent of a certain type.

Table 287: show pim statistics Output Fields *(continued)*

Field Name	Field Description
Rx errors	Number of received packets that contained errors.
V2 Hello	PIM version 2 hello packets.
V2 Register	PIM version 2 register packets.
V2 Register Stop	PIM version 2 register stop packets.
V2 Join Prune	PIM version 2 join and prune packets.
V2 Bootstrap	PIM version 2 bootstrap packets.
V2 Assert	PIM version 2 assert packets.
V2 Graft	PIM version 2 graft packets.
V2 Graft Ack	PIM version 2 graft acknowledgement packets.
V2 Candidate RP	PIM version 2 candidate RP packets.
V1 Query	PIM version 1 query packets.
V1 Register	PIM version 1 register packets.
V1 Register Stop	PIM version 1 register stop packets.
V1 Join Prune	PIM version 1 join and prune packets.
V1 RP Reachability	PIM version 1 RP reachability packets.
V1 Assert	PIM version 1 assert packets.
V1 Graft	PIM version 1 graft packets.
V1 Graft Ack	PIM version 1 graft acknowledgement packets.
AutoRP Announce	Auto-RP announce packets.
AutoRP Mapping	Auto-RP mapping packets.
AutoRP Unknown type	Auto-RP packets with an unknown type.
Anycast Register	Auto-RP announce packets.
Anycast Register Stop	Auto-RP announce packets.
Global Statistics	Summary of PIM statistics for all interfaces.

Table 287: show pim statistics Output Fields *(continued)*

Field Name	Field Description
Hello dropped on neighbor policy	Number of hello packets dropped because of a configured neighbor policy.
Unknown type	Number of PIM control packets received with an unknown type.
V1 Unknown type	Number of PIM version 1 control packets received with an unknown type.
Unknown Version	Number of PIM control packets received with an unknown version. The version is not version 1 or version 2.
Neighbor unknown	Number of PIM control packets received (excluding PIM hello) without first receiving the hello packet.
Bad Length	Number of PIM control packets received for which the packet size does not match the PIM length field in the packet.
Bad Checksum	Number of PIM control packets received for which the calculated checksum does not match the checksum field in the packet.
Bad Receive If	Number of PIM control packets received on an interface that does not have PIM configured.
Rx Bad Data	Number of PIM control packets received that contain data for TCP. Bad register packets.
Rx Intf disabled	Number of PIM control packets received on an interface that has PIM disabled.
Rx V1 Require V2	Number of PIM version 1 control packets received on an interface configured for PIM version 2.
Rx V2 Require V1	Number of PIM version 2 control packets received on an interface configured for PIM version 1.
Rx Register not RP	Number of PIM register packets received when the router is not the RP for the group.
Rx Register no route	Number of PIM register packets received when the RP does not have a unicast route back to source.
Rx Register no decap if	Number of PIM register packets received when the RP does not have a de-encapsulation interface.
Null Register Timeout	Number of NULL register timeout packets.
RP Filtered Source	Number of PIM packets received when the router has a source address filter configured for the RP.

Table 287: show pim statistics Output Fields *(continued)*

Field Name	Field Description
Rx Unknown Reg Stop	Number of register stop messages with an unknown type.
Rx Join/Prune no state	Number of join and prune messages received for which the router has no state.
Rx Join/Prune on upstream if	Number of join and prune messages received on the interface used to reach the upstream router, toward the RP.
Rx Join/Prune messages dropped	Number of join and prune messages received and dropped.
Rx sparse join for dense group	Number of PIM sparse mode join messages received for a group that is configured for dense mode.
Rx Graft/Graft Ack no state	Number of graft and graft acknowledgement messages received for which the router has no state.
Rx Graft on upstream if	Number of graft messages received on the interface used to reach the upstream router, toward the RP.
Rx CRP not BSR	Number of BSR messages received in which the PIM message type is Candidate-RP-Advertisement, not Bootstrap.
Rx BSR when BSR	Number of BSR messages received in which the PIM message type is Bootstrap.
Rx BSR not RPF if	Number of BSR messages received on an interface that is not the RPF interface.
Rx unknown hello opt	Number of PIM hello packets received with options that Junos OS does not support.
Rx data no state	Number of PIM control packets received for which the router has no state for the data type.
Rx RP no state	Number of PIM control packets received for which the router has no state for the RP.
Rx aggregate	Number of PIM aggregate MDT packets received.
Rx malformed packet	Number of PIM control packets received with a malformed IP unicast or multicast address family.
No RP	Number of PIM control packets received with no RP address.
No register encap if	Number of PIM register packets received when the first-hop router does not have an encapsulation interface.

Table 287: show pim statistics Output Fields *(continued)*

Field Name	Field Description
No route upstream	Number of PIM control packets received when the router does not have a unicast route to the the interface used to reach the upstream router, toward the RP.
Nexthop Unusable	Number of PIM control packets with an unusable nexthop. A path can be unusable if the route is hidden or the link is down.
RP mismatch	Number of PIM control packets received for which the router has an RP mismatch.
RPF neighbor unknown	Number of PIM control packets received for which the router has an unknown RPF neighbor for the source.
Rx Joins/Prunes filtered	The number of join and prune messages filtered because of configured route filters and source address filters.
Tx Joins/Prunes filtered	The number of join and prune messages filtered because of configured route filters and source address filters.
Embedded-RP invalid addr	Number of packets received with an invalid embedded RP address in PIM join messages and other types of messages sent between routing domains.
Embedded-RP limit exceed	Number of times the limit configure with the **maximum-rps** statement is exceeded. The **maximum-rps** statement limits the number of embedded RPs created in a specific routing instance. The range is from 1 through 500. The default is 100.
Embedded-RP added	Number of packets in which the embedded RP for IPv6 is added. The following receive events trigger extraction of an IPv6 embedded RP address on the router: • Multicast Listener Discovery (MLD) report for an embedded RP multicast group address • PIM join message with an embedded RP multicast group address • Static embedded RP multicast group address associated with an interface • Packets sent to an embedded RP multicast group address received on the DR An embedded RP node discovered through these receive events is added if it does not already exist on the routing platform.
Embedded-RP removed	Number of packets in which the embedded RP for IPv6 is removed. The embedded RP is removed whenever all PIM join states using this RP are removed or the configuration changes to remove the embedded RP feature.
Rx Register msgs filtering drop	Number of register messages dropped because of a filter configured for PIM register messages.

Table 287: show pim statistics Output Fields *(continued)*

Field Name	Field Description
Tx Register msgs filtering drop	Number of register messages dropped because of a filter configured for PIM register messages.

show pim statistics

```
user@host> show pim statistics
PIM Message type        Received       Sent  Rx errors

V2 Hello                      15         32          0

V2 Register                    0        362          0

V2 Register Stop             483          0          0

V2 Join Prune                 18        518          0

V2 Bootstrap                   0          0          0

V2 Assert                      0          0          0

V2 Graft                       0          0          0

V2 Graft Ack                   0          0          0

V2 Candidate RP                0          0          0

V1 Query                       0          0          0

V1 Register                    0          0          0

V1 Register Stop               0          0          0

V1 Join Prune                  0          0          0

V1 RP Reachability             0          0          0

V1 Assert                      0          0          0

V1 Graft                       0          0          0

V1 Graft Ack                   0          0          0

AutoRP Announce                0          0          0

AutoRP Mapping                 0          0          0

AutoRP Unknown type            0

Anycast Register               0          0          0

Anycast Register Stop          0          0          0


Global Statistics
```

```
Hello dropped on neighbor policy          0
Unknown type                              0
V1 Unknown type                           0
Unknown Version                           0
Neighbor unknown                          5
Bad Length                                0
Bad Checksum                              0
Bad Receive If                            0
Rx Bad Data                               0
Rx Intf disabled                          0
Rx V1 Require V2                          0
Rx V2 Require V1                          0
Rx Register not RP                        0
Rx Register no route                      0
Rx Register no decap if                   0
Null Register Timeout                     0
RP Filtered Source                        0
Rx Unknown Reg Stop                       0
Rx Join/Prune no state                    0
Rx Join/Prune on upstream if              0
Rx Join/Prune messages dropped            0
Rx sparse join for dense group            0
Rx Graft/Graft Ack no state               0
Rx Graft on upstream if                   0
Rx CRP not BSR                            0
Rx BSR when BSR                           0
Rx BSR not RPF if                         0
Rx unknown hello opt                      0
Rx data no state                          0
Rx RP no state                            0
```

```
Rx aggregate                            0
Rx malformed packet                     0
No RP                                   0
No register encap if                    0
No route upstream                       0
Nexthop Unusable                        0
RP mismatch                             0
RPF neighbor unknown                    0
Rx Joins/Prunes filtered                0
Embedded-RP invalid addr                0
Embedded-RP limit exceed                0
Embedded-RP added                       0
Embedded-RP removed                     0
Rx Register msgs filtering drop         0
Tx Register msgs filtering drop         0
```

PART 17

Access Control

CHAPTER 81

802.1X and MAC RADIUS Authentication Overview

Security Features for J-EX Series Switches Overview

The Junos operating system (Junos OS) is a network operating system that has been hardened through the separation of control forwarding and services planes, with each function running in protected memory. The control-plane CPU is protected by rate limiting, routing policy, and firewall filters to ensure switch uptime even under severe attack. In addition, the switches fully integrate with the Juniper Networks Unified Access Control (UAC) product to provide both standards-based 802.1X port-level access and Layer 2 through Layer 4 policy enforcement based on user identity. Access port security features such as dynamic Address Resolution Protocol (ARP) inspection, DHCP snooping, and MAC limiting are controlled through a single Junos OS CLI command.

J-EX Series Ethernet Switches provide the following hardware and software security features:

Console Port—Allows use of the console port to connect to the Routing Engine through an RJ-45 cable. You then use the command-line interface (CLI) to configure the switch.

Out-of-Band Management—A dedicated management Ethernet port on the rear panel allows out-of-band management.

Software Images—All Junos OS images are signed by Juniper Networks certificate authority (CA) with public key infrastructure (PKI).

User Authentication, Authorization, and Accounting (AAA)—Features include:

- User and group accounts with password encryption and authentication.
- Access privilege levels configurable for login classes and user templates.
- RADIUS authentication, TACACS+ authentication, or both, for authenticating users who attempt to access the switch.
- Auditing of configuration changes through system logging or RADIUS/TACACS+.

802.1X Authentication—Provides network access control. Supplicants (hosts) are authenticated when they initially connect to a LAN. Authenticating supplicants before they receive an IP address from a DHCP server prevents unauthorized supplicants from gaining access to the LAN. EX Series switches support Extensible Authentication Protocol (EAP) methods, including EAP-MD5, EAP-TLS, EAP-TTLS, and EAP-PEAP.

Port Security—Access port security features include:

- DHCP snooping—Filters and blocks ingress DHCP server messages on untrusted ports; builds and maintains an IP-address/MAC-address binding database (called the DHCP snooping database).
- Dynamic ARP inspection (DAI)—Prevents ARP spoofing attacks. ARP requests and replies are compared against entries in the DHCP snooping database, and filtering decisions are made based on the results of those comparisons.
- MAC limiting—Protects against flooding of the Ethernet switching table.
- MAC move limiting—Detects MAC movement and MAC spoofing on access ports.
- Trusted DHCP server—With a DHCP server on a trusted port, protects against rogue DHCP servers sending leases.
- IP source guard—Mitigates the effects of IP address spoofing attacks on the Ethernet LAN. The source IP address in the packet sent from an untrusted access interface is validated against the source MAC address in the DHCP snooping database. The packet is allowed for further processing if the source IP address to source MAC address binding is valid; if the binding is not valid, the packet is discarded.
- DHCP option 82—Also known as the DHCP relay agent information option. Helps protect the EX Series switch against attacks such as spoofing (forging) of IP addresses and MAC addresses and DHCP IP address starvation. Option 82 provides information about the network location of a DHCP client, and the DHCP server uses this information to implement IP addresses or other parameters for the client.
- Unrestricted proxy ARP—The switch responds to all ARP messages with its own MAC address. Hosts that are connected to the switch's interfaces cannot communicate

directly with other hosts. Instead, all communications between hosts go through the switch.

- Restricted proxy ARP—The switch does not respond to an ARP request if the physical networks of the source and target of the ARP request are the same. It does not matter whether the destination host has the same IP address as the incoming interface or a different (remote) IP address. An ARP request for a broadcast address elicits no reply.

Device Security—Storm control permits the switch to monitor unknown unicast and broadcast traffic and drop packets, or shut down, or temporarily disable the interface when a specified traffic level is exceeded, thus preventing packets from proliferating and degrading the LAN. You can enable storm control on access interfaces or trunk interfaces.

Firewall Filters—Allow auditing of various types of security violations, including attempts to access the switch from unauthorized locations. Firewall filters can detect such attempts and create audit log entries when they occur. The filters can also restrict access by limiting traffic to source and destination MAC addresses, specific protocols, or, in combination with policers, to specified data rates to prevent denial of service (DoS) attacks.

Policers—Provide rate-limiting capability to control the amount of traffic that enters an interface, which acts to counter DoS attacks.

Encryption Standards—Supported standards include:

- 128-, 192-, and 256-bit Advanced Encryption Standard (AES)
- 56-bit Data Encryption Standard (DES) and 168-bit 3DES

Related Documentation

- 802.1X for J-EX Series Switches Overview on page 2253
- Firewall Filters for J-EX Series Switches Overview on page 2721
- Port Security for J-EX Series Switches Overview on page 2545
- Understanding Proxy ARP on J-EX Series Switches on page 1059
- Understanding Storm Control on J-EX Series Switches on page 2511
- Understanding the Use of Policers in Firewall Filters on page 2752

Understanding Authentication on J-EX Series Switches

You can control access to your network through a J-EX Series Switch using several different authentication methods—802.1X, MAC RADIUS, or captive portal. Authentication prevents unauthorized devices and users from gaining access to your LAN. For 802.1X and MAC RADIUS authentication, end devices must be authenticated before they receive an IP address from a DHCP server. For captive portal authentication, the switch allows the devices to get an IP address, and allows DHCP, DNS and ARP packets.

You can allow end devices to access the network without authentication by including the MAC address of the end device in the static MAC bypass list or, for captive portal, by including the MAC address of the end device in the authentication whitelist.

You can configure 802.1X, MAC RADIUS, and captive portal on the same interface and in any combination, except that you cannot configure MAC RADIUS and captive portal on an interface without also configuring 802.1X. If you configure multiple authentication methods on a single interface, the switch falls back to another method if the first method is unsuccessful. For a description of the process flow when multiple authentication methods are configured on an interface, see “Authentication Process Flow for EX Series Switches” on page 2255.

This topic covers:

A Basic Authentication Topology

Figure 45 on page 2249 illustrates a basic deployment topology for authentication on a J-EX Series switch:

Figure 45: Example Authentication Topology

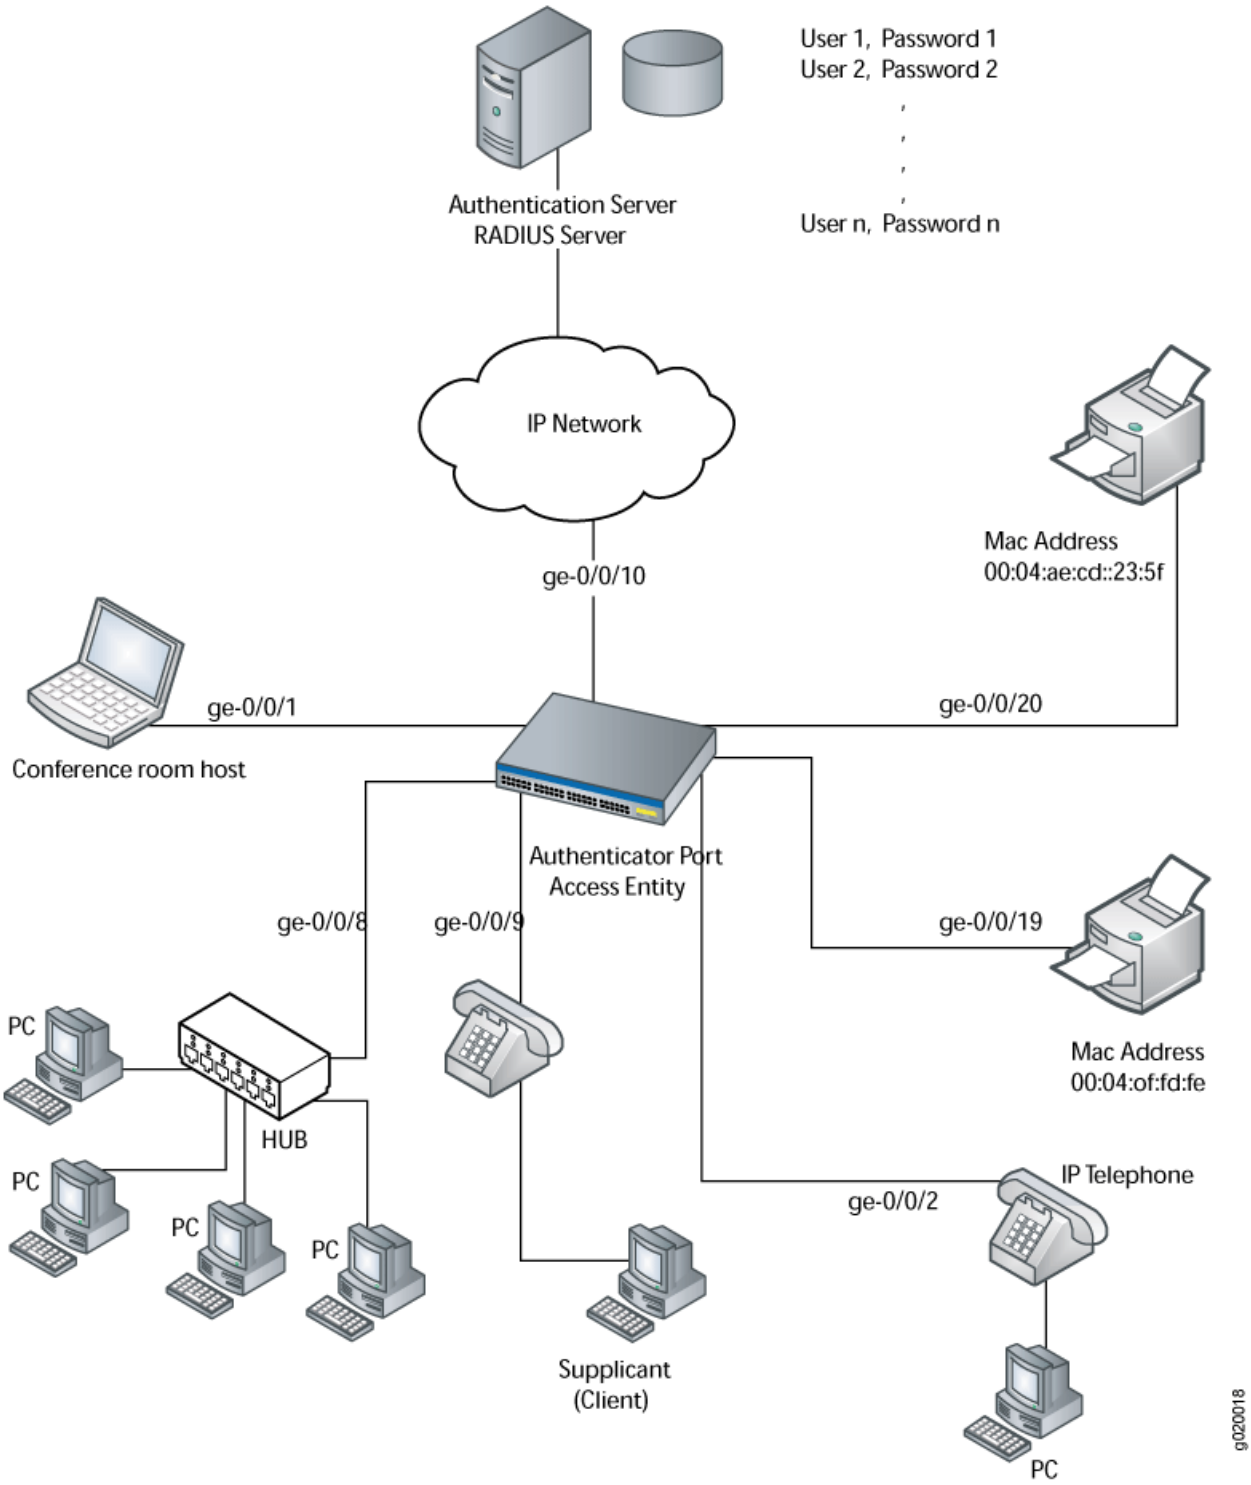

802.1X Authentication

802.1X is an IEEE standard for port-based network access control (PNAC). It provides an authentication mechanism to allow devices to access a LAN. The 802.1X authentication feature on a J-EX Series switch is based upon the IEEE 802.1D standard *Port-Based Network Access Control*.

The communication protocol between the end device and the switch is Extensible Authentication Protocol Over LAN (EAPOL). EAPOL is a version of EAP designed to work with Ethernet networks. The communication protocol between the authentication server and the switch is RADIUS.

During the authentication process, the switch completes multiple message exchanges between the end device and the authentication server. While 802.1X authentication is in process, only 802.1X traffic is allowed. Other traffic, such as DHCP and HTTP, is blocked at the data link layer.

NOTE: You can configure both the maximum number of times an EAPOL request packet is retransmitted and the timeout period between attempts. For information, see "Configuring 802.1X Interface Settings (CLI Procedure)" on page 2331.

An 802.1X authentication configuration for a LAN contains three basic components:

- *Supplicant* (also called end device)—Supplicant is the IEEE term for an end device that requests to join the network. The device can be responsive or nonresponsive. A responsive device is 802.1X-enabled and provides authentication credentials—specifically, a username and password for EAP MD5, or a username and client certificates for EAP-TLS, EAP-TTLS, and EAP-PEAP. A nonresponsive device is not 802.1X-enabled, but can be authenticated through a MAC-based authentication method.
- *Authenticator port access entity*—The IEEE term for the authenticator. The J-EX Series switch is the authenticator, and it controls access by blocking all traffic to and from end devices until they are authenticated.
- *Authentication server*—The authentication server contains the backend database that makes authentication decisions. It contains credential information for each end device that is allowed to connect to the network. The authenticator forwards credentials supplied by the end device to the authentication server. If the credentials forwarded by the authenticator match the credentials in the authentication server database, access is granted. If the credentials forwarded do not match, access is denied. The J-EX Series switches support RADIUS authentication servers.

MAC RADIUS Authentication

You can configure MAC RADIUS authentication on interfaces that are connected to end devices that are not 802.1X-enabled but that you want to allow to access the LAN.

The EAP method supported for MAC RADIUS authentication on J-EX Series switches is EAP-MD5.

If both 802.1X-enabled end-devices and end devices that are not 802.1X-enabled connect to an interface, you can configure both 802.1X and MAC RADIUS authentication methods on the interface. In this case, the switch will first attempt to authenticate using 802.1X, and if that method fails, it will attempt to authenticate the end device using MAC RADIUS authentication.

If you know that only non-802.1X-enabled end devices will connect on that interface, you can eliminate the delay that occurs while the switch determines that the end device is not non-802.1X-enabled by configuring the **mac-radius restrict** option. When this option is configured, the switch will not attempt to authenticate the end device through 802.1X but instead immediately sends a request to the RADIUS server for authentication of the MAC address of the end device. If the MAC address of an end device is configured as permitted on the RADIUS server, the switch opens LAN access to the end device on the interface to which it is connected.

This option is useful when no other 802.1X authentication methods, such as guest VLAN, are needed on the interface. When you configure **mac-radius restrict** on an interface to eliminate this delay, the switch drops all 802.1X packets.

Captive Portal Authentication

Captive portal authentication (hereafter referred to as captive portal) allows you to authenticate users on J-EX Series switches by redirecting Web browser requests to a login page that requires users to input a username and password before they are allowed access to the network. Captive portal controls network access by requiring users to provide information that is authenticated against a RADIUS server database using EAP-MD5. You can also use captive portal to display an acceptable-use policy to users before they access your network.

Junos OS for J-EX Series switches provides a template that allows you to easily design and modify the look of the captive portal login page. You enable specific interfaces for captive portal. The first time an end device connected to a captive portal interface attempts to access a web page, the switch presents the captive portal login page. Upon successful authentication, the user is allowed access to the network and to continue to the original page requested.

NOTE: If Hypertext Transfer Protocol over Secure Sockets Layer (HTTPS) is enabled, Hypertext Transfer Protocol (HTTP) requests are redirected to an HTTPS connection for the captive portal authentication process. After authentication, the end device is returned to the HTTP connection.

If there are end devices that are not HTTP-enabled connected to the captive portal interface, you can allow them to bypass captive portal authentication by adding their MAC address to an authentication whitelist.

When the user is authenticated by the RADIUS server, any per-user policies (attributes) associated with that user are also sent to the switch.

Captive portal on J-EX Series switches has the following limitations:

- The captive portal interface must be configured for **family ethernet-switching** and set to port mode **access**.
- Captive portal does not support dynamic assignment of VLANs downloaded from the RADIUS server.
- If the user is idle for more than about 5 minutes and there is no traffic passed, the user must log back in to the captive portal.

Static MAC Bypass of Authentication

You can allow end devices to access the LAN without authentication on a RADIUS server by including their MAC addresses in the static MAC bypass list (also known as the exclusion list).

You might choose to include a device in the bypass list to:

- Allow non-802.1X-enabled devices access to the LAN.
- Eliminate the delay that occurs while the switch determines that a connected device is a non-802.1X-enabled host.

When you configure static MAC on the switch, the MAC address of the end device is first checked in a local database (a user-configured list of MAC addresses). If a match is found, the end device is successfully authenticated and the interface is opened up for it. No further authentication is done for that end device. If a match is not found and 802.1X authentication is enabled on the switch, the switch attempts to authenticate the end device through the RADIUS server.

For each MAC address, you can also configure the VLAN to which the end device is moved or the interfaces on which the host connects.

Fallback of Authentication Methods

You can configure one or more authentication methods on a single interface and thus enable fallback to the next method if the first or second method fails.

If an interface is configured in multiple supplicant mode, all end devices connecting through the interface must use either captive portal or a combination of 802.1X and MAC RADIUS, captive portal cannot be mixed with 802.1X or MAC RADIUS. Therefore, if there is already an end device on the interface that was authenticated through 802.1X or MAC RADIUS authentication, then additional end devices authenticating do not fall back to captive portal. If only 802.1X authentication or MAC RADIUS authentication is configured, some end devices can be authenticated using 802.1X and others can still be authenticated using MAC RADIUS.

Fallback of authentication methods occurs in the following order:

1. 802.1X authentication—If 802.1X is configured on the interface, the switch sends EAPOL requests to the end device and attempts to authenticate the end device through 802.1X authentication. If the end device does not respond to the EAP requests, the switch checks whether MAC RADIUS authentication is configured on the interface.

2. MAC RADIUS authentication—If MAC RADIUS authentication is configured on the interface, the switch sends the MAC RADIUS address of the end device to the authentication server. If MAC RADIUS authentication is not configured, the switch checks whether captive portal is configured on the interface.
3. Captive portal authentication—If captive portal is configured on the interface, the switch attempts to authenticate using this method after attempting any other configured authentication methods. If an end device is authenticated on the interface using captive portal, this becomes the active authentication method on the interface. When captive portal is the active authentication method, the switch falls back to 802.1X authentication if there are no sessions in the authenticated state and if the interface receives an EAP packet.

Related Documentation

- 802.1X for J-EX Series Switches Overview on page 2253
- Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on a J-EX Series Switch on page 2290
- Configuring 802.1X Interface Settings (CLI Procedure) on page 2331
- Configuring MAC RADIUS Authentication (CLI Procedure) on page 2335
- Configuring MAC RADIUS Authentication (CLI Procedure) on page 2335
- Configuring Captive Portal Authentication (CLI Procedure) on page 2350
- Configuring Static MAC Bypass of Authentication (CLI Procedure) on page 2334
- Authentication Process Flow for J-EX Series Switches on page 2255

802.1X for J-EX Series Switches Overview

IEEE 802.1X provides network edge security, protecting Ethernet LANs from unauthorized user access.

How 802.1X Authentication Works

802.1X authentication works by using an *Authenticator Port Access Entity* (the switch) to block all traffic to and from a supplicant (end device) at the port until the supplicant's credentials are presented and matched on the *Authentication server* (a RADIUS server). When authenticated, the switch stops blocking traffic and opens the port to the supplicant.

The end device is authenticated in either *single* mode, *single-secure* mode, or *multiple* mode:

- **single**—Authenticates only the first end device. All other end devices that connect later to the port are allowed full access without any further authentication. They effectively “piggyback” on the end devices’ authentication.
- **single-secure**—Allows only one end device to connect to the port. No other end device is allowed to connect until the first logs out.
- **multiple**—Allows multiple end devices to connect to the port. Each end device will be authenticated individually.

Network access can be further defined using VLANs and firewall filters, which both act as filters to separate and match groups of end devices to the areas of the LAN they require.

802.1X Features Overview

802.1X features on J-EX Series Switches are:

- Guest VLAN—Provides limited access to a LAN, typically just to the Internet, for supplicants that fail 802.1X authentication.
- Server-reject VLAN—Provides limited access to a LAN, typically just to the Internet, for end devices that fail MAC RADIUS authentication.
- Dynamic VLAN—Enables a supplicant, after authentication, to be a member of a VLAN dynamically.
- Private VLAN—Enables configuration of 802.1X authentication on interfaces that are members of private VLANs (PVLANs).
- Dynamic changes to a user session—Allows the switch administrator to terminate an already authenticated session. This feature is based on support of the RADIUS Disconnect Message defined in RFC 3576.
- Support for VoIP—Supports IP telephones. If the phone is 802.1X-enabled, it is authenticated like any other supplicant. If the phone is not 802.1X-enabled, but has another 802.1X-compatible device connected to its data port, that device is authenticated, and then VoIP traffic can flow to and from the phone (providing that the interface is configured in single mode and not in single-secure mode).

  NOTE: Configuring a VoIP VLAN on private VLAN (PVLAN) interfaces is not supported.

- RADIUS accounting—Sends accounting information to the RADIUS accounting server. Accounting information is sent to the server whenever a subscriber logs in or logs out and whenever a subscriber activates or deactivates a subscription.
- Vendor Specific Attributes (VSAs)—Supports the **Juniper-Switching-Filter** attribute on the RADIUS authentication server that can be used further define a supplicant's access during the 802.1X authentication process. Centrally configuring VSAs on the authentication server does away with the need to configure these same attributes in the form of firewall filters on every switch in the LAN to which the supplicant may connect to the LAN. This feature is based on RLI 4583, AAA RADIUS BRAS VSA Support.

Supported Features Related to 802.1X Authentication

802.1X does not replace other security technologies. 802.1X works together with port security features, such as DHCP snooping, dynamic ARP inspection (DAI), and MAC limiting, to guard against spoofing.

Supported features related to authentication include:

- Static MAC bypass—Provides a bypass mechanism to authenticate devices that are not 802.1X-enabled (such as printers). Static MAC bypass connects these devices to 802.1X-enabled ports, bypassing 802.1X authentication.
- MAC RADIUS authentication—Provides a means to enable or disable MAC authentication independently of whether 802.1X authentication is enabled.

Related Documentation

- Understanding Authentication on J-EX Series Switches on page 2248
- Understanding 802.1X and VoIP on J-EX Series Switches on page 2263
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261
- Understanding 802.1X and RADIUS Accounting on J-EX Series Switches on page 2260
- Understanding Guest VLANs for 802.1X on J-EX Series Switches on page 2259
- Understanding 802.1X and VSAs on J-EX Series Switches on page 2266
- Understanding Server Fail Fallback and Authentication on J-EX Series Switches on page 2258

Authentication Process Flow for EX Series Switches

You can control access to your network through a J-EX Series switch by using several different authentication methods—802.1X, MAC RADIUS, or captive portal.

Figure 46 on page 2256 illustrates the authentication process:

Figure 46: Authentication Process Flow for a J-EX Series Switch

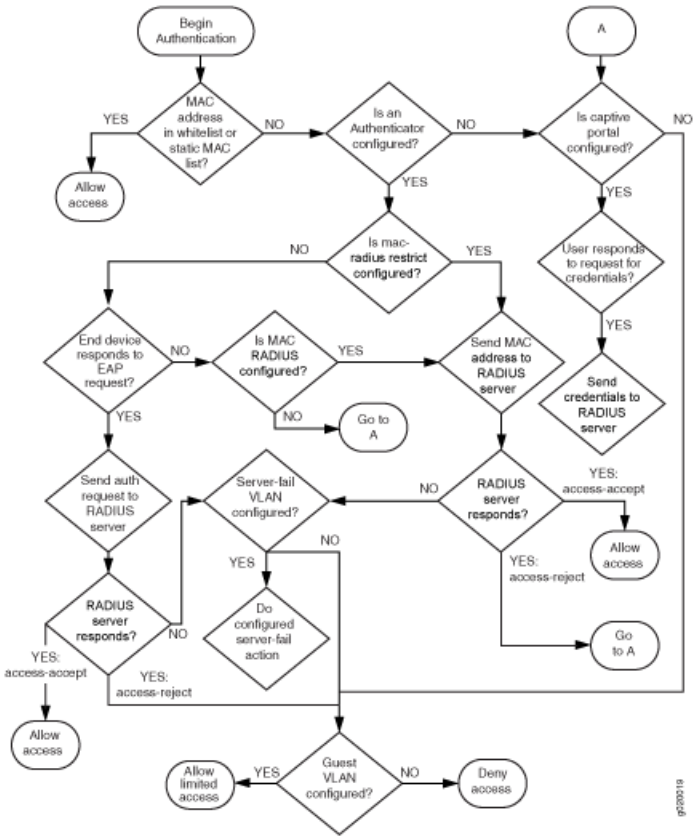

The basic authentication process works like this:

1. Authentication is initiated by an end device sending an EAP request or a data packet.
2. If the MAC address of the end device is in the static MAC bypass list or the authentication whitelist, the switch accepts the end device without querying the authentication server and allows the end device to access the LAN.
3. If the MAC address is not in the static MAC bypass list or the authentication whitelist, the switch checks whether an **authenticator** statement is configured on the interface. If an authenticator is not configured, the switch checks for captive portal configuration—skip to Step 6.

   If an authenticator is configured:

   a. The switch checks whether the **mac-radius restrict** statement is configured on the interface. If **mac-radius restrict** is configured, the switch does not attempt 802.1X authentication—skip to Step 5. If it is configured, go on to Step 2.

   b. The switch sends either an EAP request (if the end device initiated contact with a data packet) or an EAP response (if the end device initiated contact with an EAPOL-start message).

   c. If there is no response, the switch tries sending an EAP request two more times.

NOTE: You can configure both the maximum number of times an EAPOL request packet is retransmitted and the timeout period between attempts. See “Configuring 802.1X Interface Settings (CLI Procedure)” on page 2331.

d. If the end device does not respond to the EAP messages sent by the switch, the switch checks for MAC RADIUS configuration—skip to Step 4. If it does respond, go on to step 5.

e. When an EAP request is received from the end device, the switch sends an authentication request message to the authentication server.

   If the authentication server does not respond, the switch checks whether there is a server fail VLAN configured. If there is a server fail VLAN, the switch performs the configured server fail fallback operation. If there is no server fail VLAN, skip to Step 6.

f. The authentication server sends an access-accept or access-reject message. If the authentication server sends an access-reject message, skip to Step 8.

4. If the end device does not respond to the EAP messages, the switch checks whether MAC RADIUS authentication is configured on the interface. If it is not configured, skip to Step 6.
5. If MAC RADIUS authentication is configured on the interface:

   a. The switch sends a MAC RADIUS authentication request to the authentication server. The switch sends only one such request.

      If the authentication server does not respond, the switch checks whether there is a server fail VLAN configured on the switch. If there is a server fail VLAN, the switch performs the configured server fail fallback operation. If there is no server fail VLAN, skip to Step 8.

   b. The authentication server sends an access-accept or access-reject message. If the authentication server sends an access-reject message, go on to Step 6.

6. If MAC RADIUS authentication is not configured on the interface or if the authentication server responds with an access-reject message for MAC RADIUS authentication, the switch checks whether captive portal is configured on the interface. If captive portal is not configured on the interface, skip to Step 8.
7. If captive portal authentication is configured on the interface:

   a. The switch sends a request to the user on the end device for captive portal authentication information.

   b. The switch sends the captive portal authentication information to the authentication server.

   c. The authentication server sends an access-accept or access-reject message.

      If the server sends an access-reject message, go on to Step 8.

NOTE: If an end device is authenticated on the interface using captive portal, this becomes the active authentication method on the interface. When captive portal is the active authentication method, the switch falls back to 802.1X authentication if there are no sessions in the authenticated state and if the interface receives an EAP packet.

8. The switch checks whether there is a guest VLAN configured on the switch. If a guest VLAN is configured, the switch allows the end device limited access to the LAN.

Related Documentation

- Configuring Server Fail Fallback (CLI Procedure) on page 2337
- Understanding Server Fail Fallback and Authentication on J-EX Series Switches on page 2258
- Understanding Guest VLANs for 802.1X on J-EX Series Switches on page 2259
- Understanding Authentication onJ-EX Series Switches on page 2248
- Understanding Dynamic VLANs for 802.1X on J-EX Series Switches on page 2259

Understanding Server Fail Fallback and Authentication on J-EX Series Switches

Server fail fallback allows you to specify how end devices connected to the switch are supported if the RADIUS authentication server becomes unavailable or sends an Extensible Authentication Protocol Over LAN (EAPOL) access-reject message.

J-EX Series Switches use authentication to implement access control in an enterprise network. If 802.1X, MAC RADIUS, or captive portal authentication are configured on the interface, end devices are evaluated at the initial connection by an authentication (RADIUS) server. If the end device is configured on the authentication server, the device is granted access to the LAN and the J-EX Series switch opens the interface to permit access.

A RADIUS server timeout occurs if no RADIUS authentication servers are reachable when a supplicant logs in and attempts to access the LAN. Server fail fallback allows you to specify one of four actions to be taken towards end devices awaiting authentication when the server is timed out:

- *Permit* authentication, allowing traffic to flow from the end device through the interface as if the end device were successfully authenticated by the RADIUS server.
- *Deny* authentication, preventing traffic from flowing from the end device through the interface. This is the default.
- *Move* the end device to a specified VLAN. (The VLAN must already exist on the switch.)
- *Sustain* authenticated end devices that already have LAN access and *deny* unauthenticated end devices. If the RADIUS servers time out during reauthentication, previously authenticated end devices are reauthenticated and new users are denied LAN access.

Server fail fallback is triggered most often during reauthentication when the already configured and in-use RADIUS server becomes inaccessible. However, server fail fallback can also be triggered by an end device's first attempt at authentication through the RADIUS server.

Server fail fallback allows you to specify that an end device be moved to a specified VLAN if the switch receives an EAPOL accept-reject message. The configured VLAN name overrides any attributes sent by the server.

Related Documentation

- 802.1X for J-EX Series Switches Overview on page 2253
- Example: Configuring 802.1X Authentication Options When the RADIUS Server is Unavailable to a J-EX Series Switch on page 2271
- Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on a J-EX Series Switch on page 2290
- Configuring Server Fail Fallback (CLI Procedure) on page 2337
- Configuring 802.1X Interface Settings (CLI Procedure) on page 2331

Understanding Dynamic VLANs for 802.1X on J-EX Series Switches

Dynamic VLANs, in conjunction with the 802.1X authentication process, provide secure access to the LAN for supplicants belonging to different VLANs on a single port.

When this feature is configured, a supplicant becomes a member of a VLAN dynamically after 802.1X authentication is successful. Successful authentication requires that the VLAN ID or VLAN name exist on the switch and match the VLAN ID or VLAN name sent by the RADIUS server during authentication.

If the VLAN does not exist, the supplicant is unauthenticated. If a guest VLAN is established, the unauthenticated supplicant is automatically moved to the guest VLAN.

Related Documentation

- Example: Setting Up 802.1X in Conference Rooms to Provide Internet Access to Corporate Visitors on a J-EX Series Switch on page 2276
- Understanding Guest VLANs for 802.1X on J-EX Series Switches on page 2259

Understanding Guest VLANs for 802.1X on J-EX Series Switches

Guest VLANs, in conjunction with 802.1X authentication, provide secure access to the LAN for corporate guests and for supplicants who fail the 802.1X authentication process.

When a corporate visitor attempts to authenticate on the LAN, and authentication fails, the visitor is moved to a guest VLAN. A guest VLAN typically provides access only to the Internet.

A guest VLAN can also provide limited access to the LAN in cases when authentication fails for supplicants that are not visitors. When authentication fails, the switch receives an Access-Reject message for the client, and checks if a guest VLAN is configured on that port. If so, it moves that user alone to the guest VLAN. If the Access-reject message

contains optional VLAN information, then the user is moved to the VLAN specified by the RADIUS server and not to the locally configured guest-VLAN.

Authentication can fail for many reasons:

- The host device does not have supplicant software on it (for example, the host is not 802.1X-enabled, such as a printer).
- The supplicant provided invalid credentials—a username or password that were not authenticated by the authentication server.

For hosts that are not 802.1X-enabled, the guest VLAN could allow limited access to a server from which the non-802.1X-enabled host can download the supplicant software and attempt authentication again.

Related Documentation

- Example: Setting Up 802.1X in Conference Rooms to Provide Internet Access to Corporate Visitors on a J-EX Series Switch on page 2276
- Understanding Dynamic VLANs for 802.1X on J-EX Series Switches on page 2259

Understanding 802.1X and RADIUS Accounting on J-EX Series Switches

J-EX Series Switches support IETF RFC 2866, *RADIUS Accounting*. Configuring RADIUS accounting on a J-EX Series switch permits statistical data about users logging onto or off a LAN to be collected and sent to a RADIUS accounting server. The statistical data gathered can be used for general network monitoring, to analyze and track usage patterns, or to bill a user based upon the amount of time or type of services accessed.

To configure RADIUS accounting, specify one or more RADIUS accounting servers to receive the statistical data from the switch, and select the type of accounting data to be collected.

The RADIUS accounting server you specify can be the same server used for RADIUS authentication, or it can be a separate RADIUS server. You can specify a list of RADIUS accounting servers. In the event that the primary server (the first one configured) is unavailable, each RADIUS server in the list is tried in the order in which they are configured in the Junos OS.

The RADIUS accounting process between a switch and a RADIUS server works like this:

1. A RADIUS accounting server listens for User Datagram Protocol (UDP) packets on a specific port. For example, on FreeRADIUS, the default port is 1813.
2. The switch forwards an accounting-request packet containing an event record to the accounting server. For example, a supplicant is authenticated through 802.1X authentication and connected to the LAN. The event record associated with this supplicant contains an Acct-Status-Type attribute whose value indicates the beginning of user service for this supplicant. When the supplicant's session ends, the accounting request will contain an Acct-Status-Type attribute value indicating the end of user service. The RADIUS accounting server records this as a stop-accounting record containing session information and the length of the session.

3. The RADIUS accounting server logs these events as start-accounting or stop-accounting records. The records are in a file. On FreeRADIUS, the file name is the server's address; for example, 122.69.1.250.
4. The accounting server sends an accounting-response packet back to the switch confirming it has received the accounting request.
5. If the switch does not receive a response from the server, it continues to send accounting requests until an accounting response is returned from the accounting server.

The statistics collected through this process can be displayed from the RADIUS server; to see those statistics, the user accesses the log file configured to receive them.

Related Documentation

- Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch on page 2267
- 802.1X for J-EX Series Switches Overview on page 2253
- Configuring 802.1X RADIUS Accounting (CLI Procedure) on page 2339

Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches

J-EX Series Switches use Link Layer Discovery Protocol (LLDP) and Link Layer Discovery Protocol–Media Endpoint Discovery (LLDP-MED) to learn and distribute device information on network links. The information allows the switch to quickly identify a variety of devices, resulting in a LAN that interoperates smoothly and efficiently.

LLDP-capable devices transmit information in type, length, and value (TLV) messages to neighbor devices. Device information can include specifics, such as chassis and port identification and system name and system capabilities. The TLVs leverage this information from parameters that have already been configured in the Junos OS.

LLDP-MED goes one step further, exchanging IP-telephony messages between the switch and the IP telephone. These TLV messages provide detailed information on PoE policy. The PoE Management TLVs let the switch ports advertise the power level and power priority needed.

The switch also uses these protocols to ensure that voice traffic gets tagged and prioritized with the correct values at the source itself. For example, 802.1p CoS and 802.1Q tag information can be sent to the IP telephone.

J-EX Series switches support the following basic TLVs:

- **Chassis Identifier**—The MAC address associated with the local system.
- **Port identifier**—The port identification for the specified port in the local system.
- **Port Description**—The user-configured port description. The port description can be a maximum of 256 characters.
- **System Name**—The user-configured name of the local system. The system name can be a maximum of 256 characters.

- **System Description**—The system description containing information about the software and current image running on the system. This information is not configurable, but taken from the software.
- **System Capabilities**—The primary function performed by the system. The capabilities that system supports; for example, bridge or router. This information is not configurable, but based on the model of the product.
- **Management Address**—The IP management address of the local system.

J-EX Series switches support the following 802.3 TLVs:

- **Power via MDI**—A TLV that advertises MDI power support, PSE power pair, and power class information.
- **MAC/PHY Configuration Status**—A TLV that advertises information about the physical interface, such as autonegotiation status and support and MAU type. The information is not configurable, but based on the physical interface structure.
- **Link Aggregation**—A TLV that advertises if the port is aggregated and its aggregated port ID.
- **Maximum Frame Size**—A TLV that advertises the Maximum Transmission Unit (MTU) of the interface sending LLDP frames.
- **Port Vlan**—A TLV that advertises the VLAN name configured on the interface.

J-EX Series switches support the following LLDP-MED TLVs:

- **LLDP MED Capabilities**—A TLV that advertises the primary function of the port. The capabilities values range 0 through 15:
  - **0**—Capabilities
  - **1**—Network Policy
  - **2**—Location Identification
  - **3**—Extended Power via MDI-PSE
  - **4**—Inventory
  - **5–15**—Reserved
- LLDP-MED Device Class Values:
  - **0**—Class not defined.
  - **1**—Class 1 Device.
  - **2**—Class 2 Device.
  - **3**—Class 3 Device.
  - **4**—Network Connectivity Device
  - **5–255**—Reserved.

- **Network Policy**—A TLV that advertises the port VLAN configuration and associated Layer 2 and Layer 3 attributes. Attributes include the policy identifier, application types, such as voice or streaming video, 802.1Q VLAN tagging, and 802.1p priority bits and Diffserv code points.
- **Endpoint Location**—A TLV that advertises the physical location of the endpoint.
- **Extended Power via MDI**—A TLV that advertises the power type, power source, power priority, and power value of the port. It is the responsibility of the PSE device (network connectivity device) to advertise the power priority on a port.

Related Documentation

- Understanding Layer 2 Protocol Tunneling on J-EX Series Switches on page 1056
- Example: Setting Up VoIP with 802.1X and LLDP-MED on a J-EX Series Switch on page 2302
- Configuring LLDP-MED (CLI Procedure) on page 2346
- Configuring LLDP (CLI Procedure) on page 2344

Understanding 802.1X and VoIP on J-EX Series Switches

When you use Voice over IP (VoIP), you can connect IP telephones to the switch and configure IEEE 802.1X authentication for 802.1X-compatible IP telephones. The 802.1X authentication provides network edge security, protecting Ethernet LANs from unauthorized user access.

VoIP is a protocol used for the transmission of voice through packet-switched networks. VoIP transmits voice calls using a network connection instead of an analog phone line.

When VoIP is used with 802.1X, the RADIUS server authenticates the phone, and Link Layer Discovery Protocol–Media Endpoint Discovery (LLDP-MED) provides the class-of-service (CoS) parameters to the phone.

You can configure 802.1X authentication to work with VoIP in multiple supplicant or single supplicant mode. In *multiple-supplicant* mode, the 802.1X process allows multiple supplicants to connect to the interface. Each supplicant will be authenticated individually. For an example of a VoIP multiple supplicant topology, see Figure 47 on page 2264.

Figure 47: VoIP Multiple Supplicant Topology

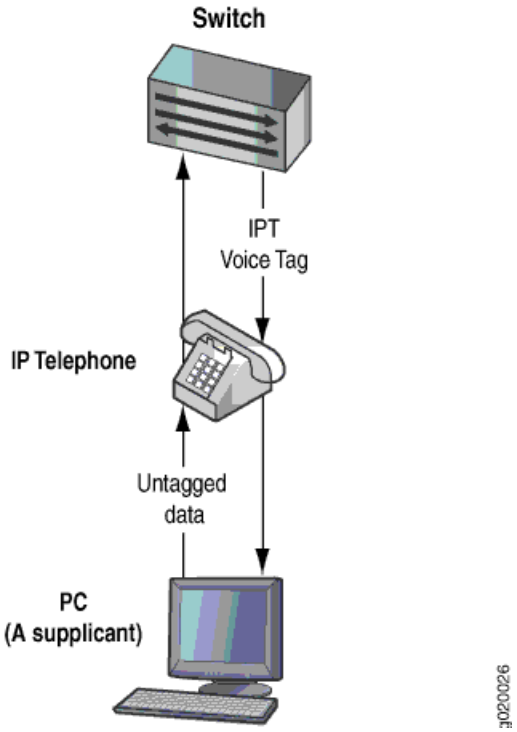

If an 802.1X-compatible IP telephone does not have an 802.1X host but has another 802.1X-compatible device connected to its data port, you can connect the phone to an interface in single-supplicant mode. In *single-supplicant* mode, the 802.1X process authenticates only the first supplicant. All other supplicants who connect later to the interface are allowed full access without any further authentication. They effectively "piggyback" on the first supplicant's authentication. For an example of a VoIP single supplicant topology, see Figure 48 on page 2265 .

Figure 48: VoIP Single Supplicant Topology

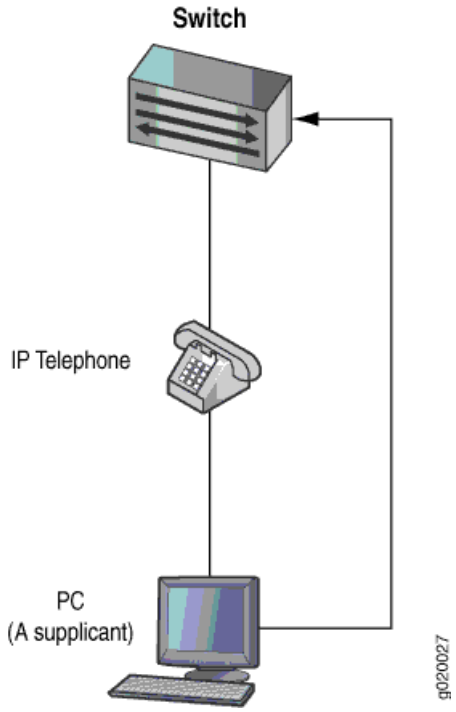

If an IP telephone does not support 802.1X, you can configure VoIP to bypass 802.1X and LLDP-MED and have the packets forwarded to a VoIP VLAN,

Related Documentation

- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261
- Example: Setting Up VoIP with 802.1X and LLDP-MED on a J-EX Series Switch on page 2302
- Example: Configuring VoIP on a J-EX Series Switch Without Including 802.1X Authentication on page 2309
- Example: Configuring VoIP on a J-EX Series Switch Without Including LLDP-MED Support on page 2315

Understanding 802.1X and VSAs on J-EX Series Switches

J-EX Series Switches support the configuration of RADIUS server attributes specific to Juniper Networks. These attributes are known as vendor-specific attributes (VSAs) and are described in RFC 2138, *Remote Authentication Dial In User Service* (RADIUS). Through VSAs, you can configure port-filtering attributes on the RADIUS server. VSAs are clear text fields sent from the RADIUS server to the switch as a result of the 802.1X authentication success or failure. The 802.1X authentication prevents unauthorized user access by blocking a supplicant at the port until the supplicant is authenticated by the RADIUS server. The VSA attributes are interpreted by the switch during authentication, and the switch takes appropriate actions. Implementing port-filtering attributes with 802.1X authentication on the RADIUS server provides a central location for controlling LAN access for supplicants.

These port-filtering attributes specific to Juniper Networks are encapsulated in a RADIUS server VSA with the vendor ID set to the Juniper Networks ID number, 2636.

As well as configuring port-filtering attributes through VSAs, you can apply a port firewall filter that has already been configured on the switch directly to the RADIUS server. Like port-filtering attributes, the filter is applied during the 802.1X authentication process, and its actions are applied at the switch port. Adding a port firewall filter to a RADIUS server eliminates the need to add the filter to multiple ports and switches. For more information, see “Example: Applying a Firewall Filter to 802.1X-Authenticated Supplicants Using RADIUS Server Attributes on a J-EX Series Switch” on page 2296.

VSAs are only supported for 802.1X single-supplicant configurations and multiple-supplicant configurations.

Related Documentation

- Understanding Authentication onJ-EX Series Switches on page 2248
- Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on a J-EX Series Switch on page 2290
- Filtering 802.1X Supplicants Using RADIUS Server Attributes on page 2340
- Configuring Firewall Filters (CLI Procedure) on page 2779
- VSA Match Conditions and Actions for J-EX Series Switches on page 2348

CHAPTER 82

Examples: Access Control Configuration

Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch

802.1X is the IEEE standard for Port-Based Network Access Control (PNAC). You use 802.1X to control network access. Only users and devices providing credentials that have been verified against a user database are allowed access to the network. You can use a RADIUS server as the user database for 802.1X authentication, as well as for MAC RADIUS authentication.

This example describes how to connect a RADIUS server to a J-EX Series switch, and configure it for 802.1X:

Requirements

This example uses the following hardware and software components:

- One J-EX Series switch acting as an authenticator port access entity (PAE). The ports on the authenticator PAE form a control gate that blocks all traffic to and from supplicants until they are authenticated.
- One RADIUS authentication server that supports 802.1X. The authentication server acts as the backend database and contains credential information for hosts (supplicants) that have permission to connect to the network.

Before you connect the server to the switch, be sure you have:

- Performed basic bridging and VLAN configuration on the switch. See “Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch” on page 1063.
- Configured users on the RADIUS authentication server.

Overview and Topology

The J-EX Series switch acts as an authenticator Port Access Entity (PAE). It blocks all traffic and acts as a control gate until the supplicant (client) is authenticated by the server. All other users and devices are denied access.

Figure 49 on page 2269 shows one J-EX4200 switch that is connected to the devices listed in Table 288 on page 2270.

Figure 49: Topology for Configuration

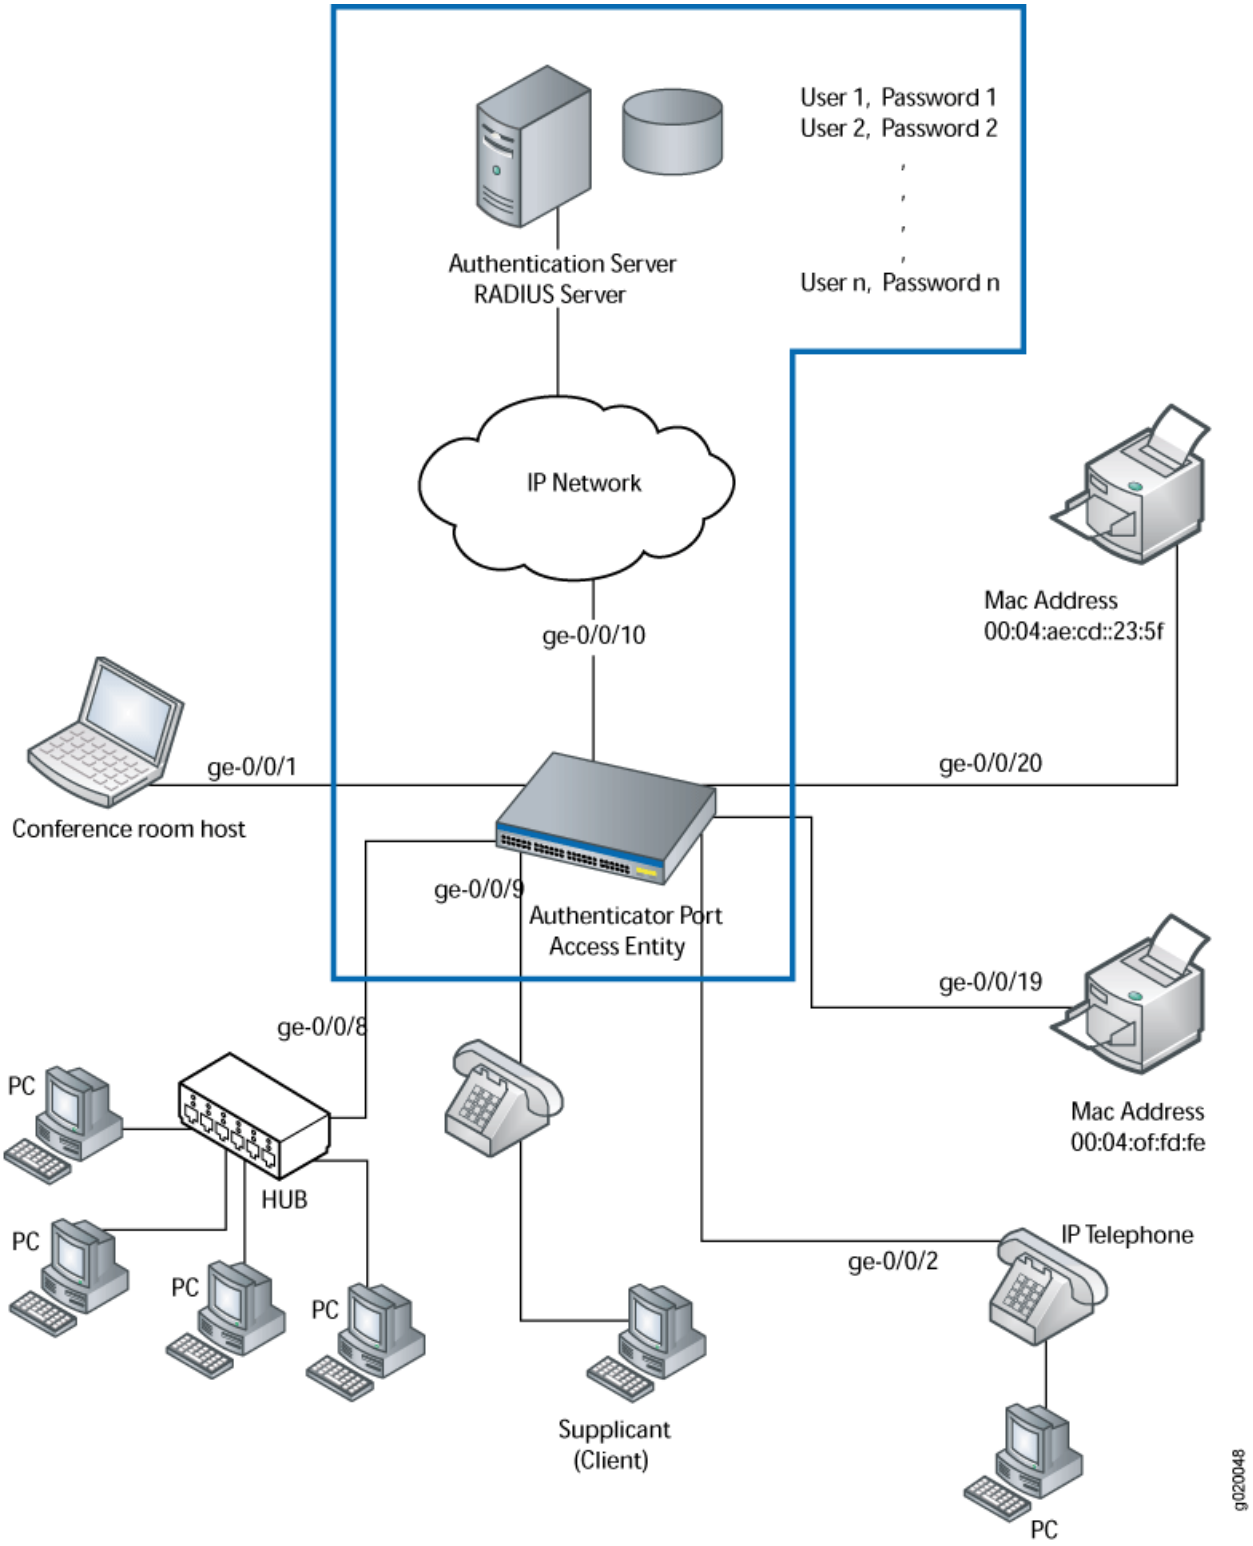

Table 288: Components of the Topology

Property	Settings
Switch hardware	J-EX4200 access switch, 24 Gigabit Ethernet ports: 8 PoE ports (**ge-0/0/0** through **ge-0/0/7**) and 16 non-PoE ports (**ge-0/0/8** through **ge-0/0/23**)
VLAN name	**default**
One RADIUS server	Backend database with an address of **10.0.0.100** connected to the switch at port **ge-0/0/10**

In this example, connect the RADIUS server to access port **ge-0/0/10** on the J-EX4200 switch. The switch acts as the authenticator and forwards credentials from the supplicant to the user database on the RADIUS server. You must configure connectivity between the J-EX4200 and the RADIUS server by specifying the address of the server and configuring the secret password. This information is configured in an access profile on the switch.

NOTE: For more information about authentication, authorization, and accounting (AAA) services, please see the *Junos OS System Basics Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.

Configuration

CLI Quick Configuration

To quickly connect the RADIUS server to the switch, copy the following commands and paste them into the switch terminal window:

```
[edit]
set access radius-server 10.0.0.100 secret juniper
set access profile profile1 authentication-order radius
set access profile profile1 radius authentication-server 10.0.0.100 10.2.14.200
```

Step-by-Step Procedure

To connect the RADIUS server to the switch:

1. Define the address of the server, and configure the secret password. The secret password on the switch must match the secret password on the server:

   ```
   [edit access]
   user@switch# set radius-server 10.0.0.100 secret juniper
   ```

2. Configure the authentication order, making **radius** the first method of authentication:

   ```
   [edit access profile]
   user@switch# set profile1 authentication-order radius
   ```

3. Configure a list of server IP addresses to be tried in order to authenticate the supplicant:

   ```
   [edit access profile]
   user@switch# set profile1 radius authentication-server 10.0.0.100 10.2.14.200
   ```

Results

Display the results of the configuration:

```
user@switch> show configuration access
```

```
radius-server {
  10.0.0.100
    port 1812;
    secret "$9$qPT3ApBSrv69rvWLVb.P5"; ## SECRET-DATA
    }
  }
  profile profile1{
    authentication-order radius;
    radius {
      authentication-server 10.0.0.100 10.2.14.200;
    }
  }
}
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verify That the Switch and RADIUS Server are Properly Connected on page 2271

Verify That the Switch and RADIUS Server are Properly Connected

Purpose Verify that the RADIUS server is connected to the switch on the specified port.

Action Ping the RADIUS server to verify the connection between the switch and the server:

```
user@switch> ping 10.0.0.100
PING 10.0.0.100 (10.0.0.100): 56 data bytes
64 bytes from 10.93.15.218: icmp_seq=0 ttl=64 time=9.734 ms
64 bytes from 10.93.15.218: icmp_seq=1 ttl=64 time=0.228 ms
```

Meaning ICMP echo request packets are sent from the switch to the target server at 10.0.0.100 to test whether it is reachable across the IP network. ICMP echo responses are being returned from the server, verifying that the switch and the server are connected.

Related Documentation

- Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on a J-EX Series Switch on page 2290
- Example: Setting Up 802.1X in Conference Rooms to Provide Internet Access to Corporate Visitors on a J-EX Series Switch on page 2276
- Example: Setting Up VoIP with 802.1X and LLDP-MED on a J-EX Series Switch on page 2302
- Configuring 802.1X RADIUS Accounting (CLI Procedure) on page 2339
- Filtering 802.1X Supplicants Using RADIUS Server Attributes on page 2340

Example: Configuring 802.1X Authentication Options When the RADIUS Server is Unavailable to a J-EX Series Switch

Server fail fallback allows you to specify how 802.1X supplicants connected to the switch are supported if the RADIUS authentication server becomes unavailable or sends an EAP Access-Reject message.

You use 802.1X to control network access. Only users and devices (supplicants) providing credentials that have been verified against a user database are allowed access to the network. You use a RADIUS server as the user database.

This example describes how to configure an interface to move a supplicant to a VLAN in the event of a RADIUS server timeout:

Requirements

This example uses the following hardware and software components:

- One J-EX Series switch acting as an authenticator port access entity (PAE). The ports on the authenticator PAE form a control gate that blocks all traffic to and from supplicants until they are authenticated.
- One RADIUS authentication server that supports 802.1X. The authentication server acts as the backend database and contains credential information for hosts (supplicants) that have permission to connect to the network.

Before you connect the server to the switch, be sure you have:

- Performed basic bridging and VLAN configuration on the switch. See “Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch” on page 1063.
- Set up a connection between the switch and the RADIUS server. See “Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch” on page 2267.
- Disable firewall filters on the interface. Firewall filters interfere with server fail fallback operation.
- Configured users on the authentication server.

Overview and Topology

A RADIUS server timeout occurs if no authentication RADIUS servers are reachable when a supplicant logs in and attempts to access the LAN. Using server fail fallback, configure alternative options for supplicants attempting LAN access. You can configure the switch to accept or deny access to supplicants or to maintain the access already granted towards supplicants before the RADIUS server timeout. Additionally, you can configure the switch to move supplicants to a specific VLAN if a RADIUS timeout occurs or if the RADIUS server sends an EAP Access-Reject message. Figure 50 on page 2273 shows the topology used for this example. The RADIUS server is connected to the J-EX4200 switch on access port **ge-0/0/10**. The switch acts as the authenticator Port Access Entity (PAE) and forwards credentials from the supplicant to the user database on the RADIUS server. The switch blocks all traffic and acts as a control gate until the supplicant is authenticated by the authentication server. A supplicant is connected to the switch through interface **ge-0/0/1**.

Figure 50: Topology for Configuration

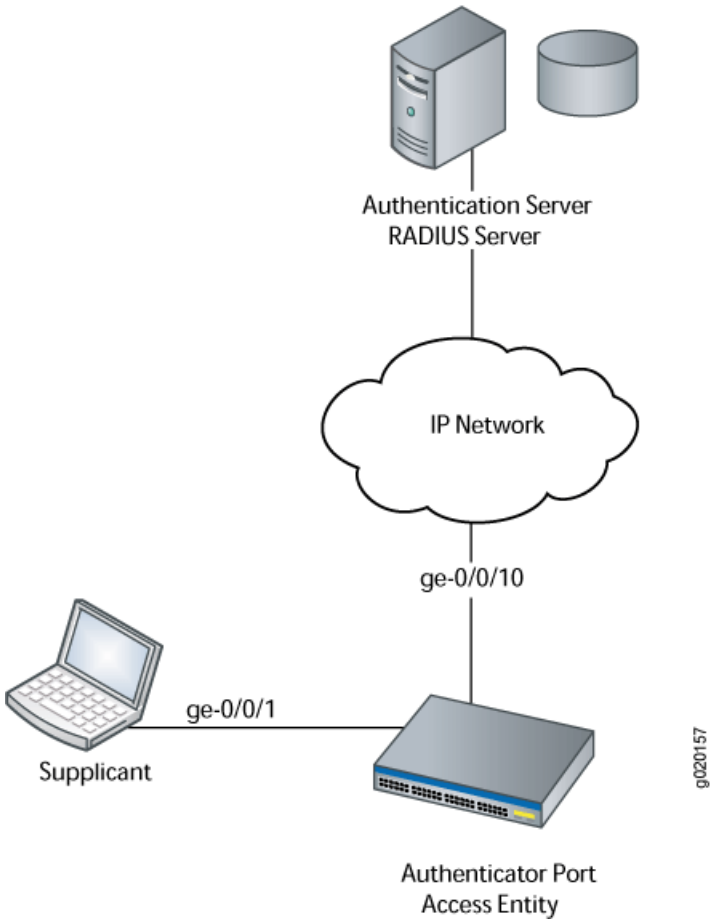

Table 289 on page 2273 describes the components in this topology.

Table 289: Components of the Topology

Property	Settings
Switch hardware	J-EX4200 access switch, 24 Gigabit Ethernet ports: 8 PoE ports.
VLAN names	**default** VLAN **vlan-sf** VLAN
Supplicant	Supplicant attempting access on interface **ge-0/0/1**
One RADIUS server	Backend database with an address of **10.0.0.100** connected to the switch at port **ge-0/0/10**

In this example, configure interface **ge-0/0/1** to move a supplicant attempting access to the LAN during a RADIUS timeout to another VLAN. A RADIUS timeout prevents the normal exchange of EAP messages that carry information from the RADIUS server to the switch and permit the authentication of a supplicant. The **default** VLAN is configured on

interface **ge-0/0/1**. When a RADIUS timeout occurs, supplicants on the interface will be moved from the **default** VLAN to the VLAN named **vlan-sf**.

NOTE: For more information about authentication, authorization, and accounting (AAA) services, see the *Junos OS System Basics Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.

Configuration

To configure server fail fallback on the switch:

CLI Quick Configuration

To quickly configure server fail fallback on the switch, copy the following commands and paste them into the switch terminal window:

```
[edit protocols dot1x authenticator]
set interface ge-0/0/1 server-fail vlan-name vlan-sf
```

Step-by-Step Procedure

To configure an interface to divert supplicants to a specific VLAN when a RADIUS timeout occurs (here, the VLAN is **vlan-sf**):

1. Define the VLAN to which supplicants are diverted:

```
[edit protocols dot1x authenticator]
user@switch# set interface server-fail vlan-name vlan-sf
```

Results

Display the results of the configuration:

```
user@switch> show configuration
  interfaces {
    ge-0/0/1 {
      unit 0 {
        family ethernet-switching {
          vlan {
            members default;
          }
        }
      }
    }
  protocols {
    dot1x {
      authenticator {
        authentication-profile-name profile52;
        interface {
          ge-0/0/1.0 {
            server-fail vlan-name vlan-sf;
          }
        }
      }
    }
  }
}
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying That the Supplicants Are Moved to an Alternative VLAN During a RADIUS Timeout on page 2275

Verifying That the Supplicants Are Moved to an Alternative VLAN During a RADIUS Timeout

Purpose Verify that the interface moves supplicants to an alternative VLAN during a RADIUS timeout.

Action Display the VLANs configured on the switch; the interface **ge-0/0/1.0** is a member of the **default** VLAN:

```
user@switch> show vlans
Name            Tag     Interfaces
default
                        ge-0/0/0.0, ge-0/0/1.0*, ge-0/0/5.0*, ge-0/0/10.0,
                        ge-0/0/12.0*, ge-0/0/14.0*, ge-0/0/15.0, ge-0/0/20.0
v2              77
                        None
vlan-sf         50
                        None
mgmt
                        me0.0*
```

Display 802.1X protocol information on the switch to view supplicants that are authenticated on interface **ge-0/0/1.0**:

```
user@switch> show dot1x interface brief
802.1X Information:
Interface     Role           State           MAC address          User
ge-0/0/1.0    Authenticator  Authenticated   00:00:00:00:00:01    abc
ge-0/0/10.0   Authenticator  Initialize
ge-0/0/14.0   Authenticator  Connecting
ge-0/0/15.0   Authenticator  Initialize
ge-0/0/20.0   Authenticator  Initialize
```

A RADIUS server timeout occurs. Display the Ethernet switching table to show that the supplicant with the MAC address **00:00:00:00:00:01** previously accessing the LAN through the **default** VLAN is now being learned on the VLAN named **vlan-sf**:

```
user@switch> show ethernet-switching table
Ethernet-switching table: 3 entries, 1 learned
  VLAN              MAC address       Type         Age Interfaces
  v1                *                 Flood          - All-members
  vlan-sf           00:00:00:00:00:01 Learn       1:07 ge-0/0/1.0
  default           *                 Flood          - All-members
```

Display 802.1X protocol information to show that interface **ge-0/0/1.0** is connecting and will open LAN access to supplicants:

```
user@switch> show dot1x interface brief
```

```
802.1X Information:
Interface     Role           State          MAC address          User
ge-0/0/1.0    Authenticator  Connecting
ge-0/0/10.0   Authenticator  Initialize
ge-0/0/14.0   Authenticator  Connecting
ge-0/0/15.0   Authenticator  Initialize
ge-0/0/20.0   Authenticator  Initialize
```

Meaning The command **show vlans** displays interface **ge-0/0/1.0** as a member of the **default** VLAN. The command **show dot1x interface brief** shows that a supplicant (**abc**) is authenticated on interface **ge-0/0/1.0** and has the MAC address **00:00:00:00:00:01**. A RADIUS server timeout occurs, and the authentication server cannot be reached by the switch. The command **show-ethernet-switching table** shows that MAC address **00:00:00:00:00:01** is learned on VLAN **vlan-sf**. The supplicant has been moved from the **default** VLAN to the **vlan-sf** VLAN. The supplicant is then connected to the LAN through the VLAN named **vlan-sf**.

Related Documentation

- Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on a J-EX Series Switch on page 2290
- Configuring Server Fail Fallback (CLI Procedure) on page 2337
- Configuring 802.1X RADIUS Accounting (CLI Procedure) on page 2339
- Filtering 802.1X Supplicants Using RADIUS Server Attributes on page 2340
- Understanding Server Fail Fallback and Authentication on J-EX Series Switches on page 2258

Example: Setting Up 802.1X in Conference Rooms to Provide Internet Access to Corporate Visitors on a J-EX Series Switch

802.1X on J-EX Series switches provides LAN access to users who do not have credentials in the RADIUS database. These users, referred to as guests, are authenticated and typically provided with access to the Internet.

This example describes how to create a guest VLAN and configure 802.1X authentication for it.

- Requirements on page 2276
- Overview and Topology on page 2277
- Configuration of a Guest VLAN That Includes 802.1X Authentication on page 2279
- Verification on page 2280

Requirements

This example uses the following hardware and software components:

- One J-EX Series switch acting as an authenticator interface access entity (PAE). The interfaces on the authenticator PAE form a control gate that blocks all traffic to and from supplicants until they are authenticated.

- One RADIUS authentication server that supports 802.1X. The authentication server acts as the backend database and contains credential information for hosts (supplicants) that have permission to connect to the network.

Before you configure guest VLAN authentication, be sure you have:

- Installed your J-EX Series switch. See Installing and Connecting a J-EX4200 Switch.
- Performed the initial switch configuration. See “Connecting and Configuring a J-EX Series Switch (J-Web Procedure)” on page 163.
- Performed basic bridging and VLAN configuration on the switch. See “Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch” on page 1063.

Overview and Topology

As part of IEEE 802.1X Port-Based Network Access Control (PNAC), you can provide limited network access to supplicants who do not belong to a VLAN authentication group by configuring authentication to a guest VLAN. Typically, guest VLAN access is used to provide Internet access to visitors to a corporate site. However, you can also use the guest VLAN feature to provide supplicants that fail 802.1X authentication to a corporate LAN with access to a VLAN with limited resources.

Figure 51 on page 2278 shows the conference room connected to the switch at interface **ge-0/0/1**.

Figure 51: Topology for Guest VLAN Example

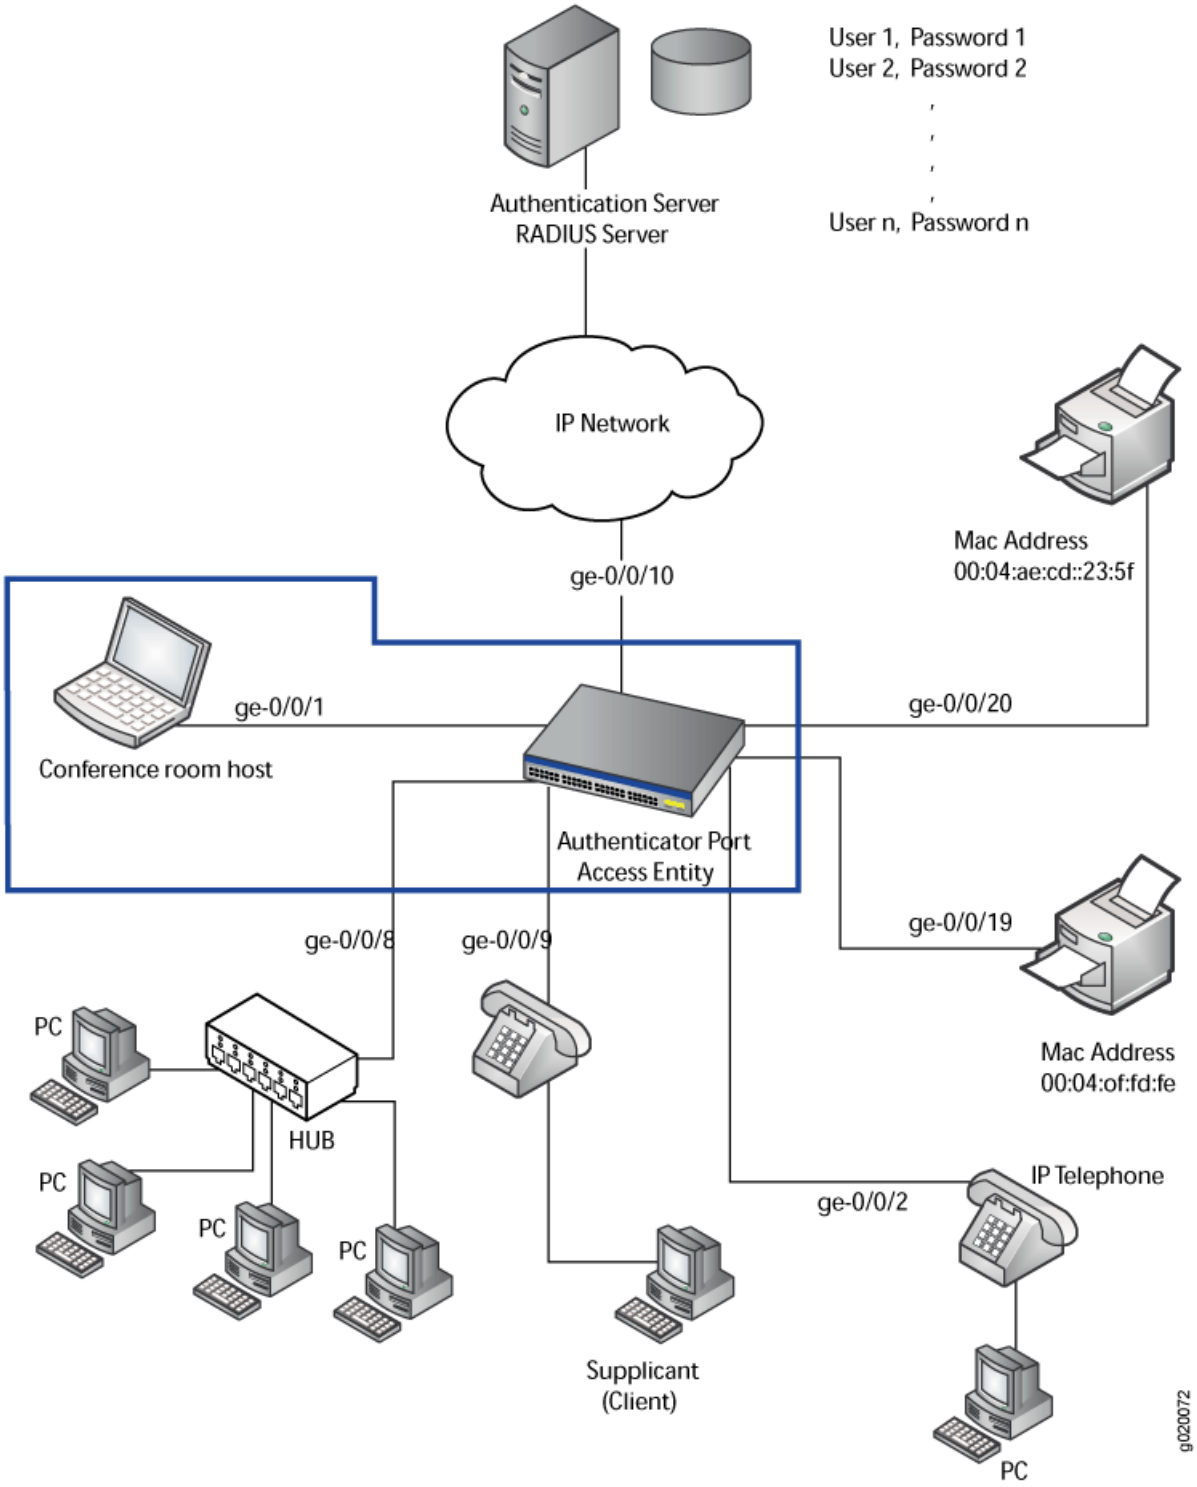

Table 290: Components of the Guest VLAN Topology

Property	Settings
Switch hardware	J-EX4200 switch, 24 Gigabit Ethernet interfaces: 8 PoE interfaces (**ge-0/0/0** through **ge-0/0/7**) and 16 non-PoE interfaces (**ge-0/0/8** through **ge-0/0/23**)
VLAN names and tag IDs	**sales**, tag **100** **support**, tag **200** **guest-vlan**, tag **300**
One RADIUS server	Backend database connected to the switch through interface **ge-0/0/10**

In this example, access interface **ge-0/0/1** provides LAN connectivity in the conference room. Configure this access interface to provide LAN connectivity to visitors in the conference room who are not authenticated by the corporate VLAN.

Configuration of a Guest VLAN That Includes 802.1X Authentication

To create a guest VLAN and configure 802.1X authentication, perform these tasks:

CLI Quick Configuration

To quickly configure a guest VLAN, with 802.1X authentication, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans guest-vlan vlan-id 300
set protocols dot1x authenticator interface all guest-vlan guest-vlan
```

Step-by-Step Procedure

To configure a guest VLAN that includes 802.1X authentication on a J-EX Series switch:

1. Configure the VLAN ID for the guest VLAN:

   ```
   [edit]

   user@switch# set vlans guest-vlan vlan-id 300
   ```

2. Configure the guest VLAN under **dot1x** protocols:

   ```
   [edit]

   user@switch# set protocols dot1x authenticator interface all guest-vlan guest-vlan
   ```

Results

Check the results of the configuration:

```
user@switch> show configuration
protocols {
    dot1x {
        authenticator {
            interface {
                all {
                    guest-vlan {
                        guest-vlan;
                    }
                }
            }
        }
    }
```

```
}
vlans {
  guest-vlan {
    vlan-id 300;
  }
}
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying That the Guest VLAN is Configured on page 2280

Verifying That the Guest VLAN is Configured

Purpose Verify that the guest VLAN is created and that an interface has failed authentication and been moved to the guest VLAN.

Action Use the operational mode commands:

```
user@switch> show vlans

Name            Tag     Interfaces
default
                        ge-0/0/3.0*
dynamic         40
                        None
guest           30
                        None
guest-vlan      300
                        ge-0/0/1.0*
vlan_dyn
                        None
```

```
user@switch> show dot1x interface ge-0/0/1.0 detail
ge-0/0/1.0
  Role: Authenticator
  Administrative state: Auto
  Supplicant mode: Single
  Number of retries: 3
  Quiet period: 60 seconds
  Transmit period: 30 seconds
  Mac Radius: Enabled
  Mac Radius Restrict: Disabled
  Reauthentication: Enabled
  Configured Reauthentication interval: 3600 seconds
  Supplicant timeout: 30 seconds
  Server timeout: 30 seconds
  Maximum EAPOL requests: 2
  Guest VLAN member: guest-vlan
  Number of connected supplicants: 1
    Supplicant: user1, 00:00:00:00:13:23
      Operational state: Authenticated
      Authentication method: Radius
      Authenticated VLAN: vo11
      Dynamic Filter: match source-dot1q-tag 10 action deny
```

```
Session Reauth interval: 60 seconds
Reauthentication due in 50 seconds
```

Meaning The output from the **show vlans** command shows **guest-vlan** as the the name of the VLAN and the VLAN ID as **300**.

The output from the **show dot1x interface ge-0/0/1.0 detail** command displays the **Guest VLAN membership** field, indicating that a supplicant at this interface failed 802.1X authentication and was passed through to the **guest-vlan**.

Related Documentation

- Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch on page 2267
- Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on a J-EX Series Switch on page 2290
- Example: Setting Up VoIP with 802.1X and LLDP-MED on a J-EX Series Switch on page 2302
- Configuring 802.1X Interface Settings (CLI Procedure) on page 2331

Example: Configuring Static MAC Bypass of Authentication on a J-EX Series Switch

To allow devices to access your LAN through 802.1X-configured interfaces without authentication, you can configure a static MAC bypass list on the J-EX Series switch. The static MAC bypass list, also known as the exclusion list, specifies MAC addresses that are allowed on the switch without a request to an authentication server.

You can use static MAC bypass of authentication to allow connection for devices that are not 802.1X-enabled, such as printers. If a host's MAC address is compared and matched against the static MAC address list, the nonresponsive host is authenticated and an interface opened for it.

This example describes how to configure static MAC bypass of authentication for two printers:

- Requirements on page 2281
- Overview and Topology on page 2282
- Configuration on page 2284
- Verification on page 2285

Requirements

This example uses the following hardware and software components:

- One J-EX Series switch acting as an authenticator port access entity (PAE). The ports on the authenticator PAE form a control gate that blocks all traffic to and from supplicants until they are authenticated.

Before you configure static MAC authentication, be sure you have:

- Performed basic bridging and VLAN configuration on the switch. See “Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch” on page 1063.

Overview and Topology

To permit printers access to the LAN, add them to the static MAC bypass list. The MAC addresses on this list are permitted access without authentication from the RADIUS server.

Figure 52 on page 2283 shows the two printers connected to the J-EX4200.

Figure 52: Topology for Static MAC Authentication Configuration

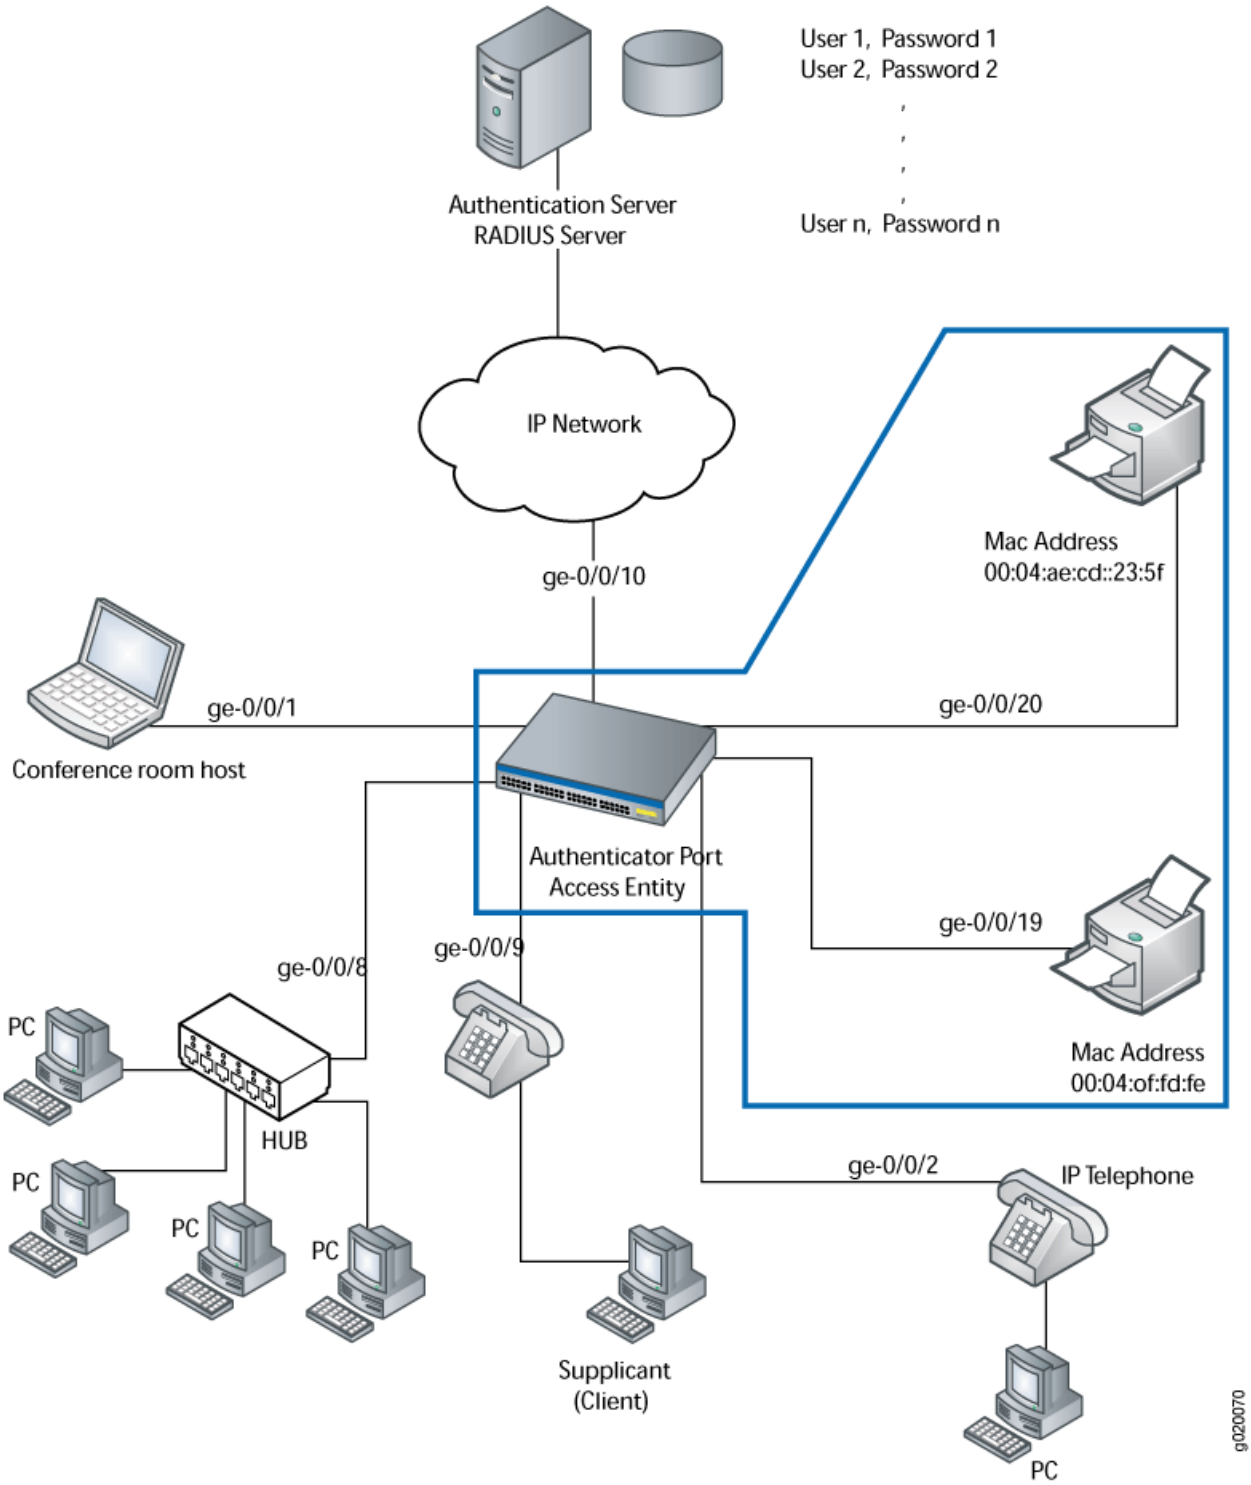

The interfaces shown in Table 291 on page 2284 will be configured for static MAC authentication.

Table 291: Components of the Static MAC Authentication Configuration Topology

Property	Settings
Switch hardware	J-EX4200, 24 Gigabit Ethernet ports: 8 PoE ports (**ge-0/0/0** through **ge-0/0/23**)
VLAN name	**default**
Connections to integrated printer/fax/copier machines (no PoE required)	**ge-0/0/19**, MAC address 00:04:0f:fd:ac:fe **ge-0/0/20**, MAC address 00:04:ae:cd:23:5f

The printer with the MAC address 00:04:0f:fd:ac:fe is connected to access interface **ge-0/0/19**. A second printer with the MAC address 00:04:ae:cd:23:5f is connected to access interface **ge-0/0/20**. Both printers will be added to the static list and bypass 802.1X authentication.

Configuration

To configure static MAC authentication, perform these tasks:

CLI Quick Configuration

To quickly configure static MAC authentication, copy the following commands and paste them into the switch terminal window:

```
[edit]
set protocols dot1x authenticator authenticaton-profile-name profile1
set protocols dot1x authenticator static [00:04:0f:fd:ac:fe 00:04:ae:cd:23:5f]
set protocols dot1x interface all supplicant multiple
```

Step-by-Step Procedure

Configure static MAC authentication:

1. Configure the authentication profile name (access profile name) to use for authentication:

   ```
   [edit protocols]
   user@switch# set dot1x authenticator authentication-profile-name profile1
   ```

2. Configure MAC addresses **00:04:0f:fd:ac:fe** and **00:04:ae:cd:23:5f** as static MAC addresses:

   ```
   [edit protocols]
   user@switch# set dot1x authenticator static [00:04:0f:fd:ac:fe 00:04:ae:cd:23:5f]
   ```

3. Configure the 802.1X authentication method:

   ```
   [edit protocols]
   user@switch# set dot1x interface all supplicant multiple
   ```

Results

Display the results of the configuration:

```
user@switch> show
interfaces {
  ge-0/0/19 {
    unit 0 {
      family ethernet-switching {
        vlan members default;
      }
    }
```

```
    }
    ge-0/0/20 {
      unit 0 {
        family ethernet-switching {
          vlan members default;
        }
      }
    }
  }
  protocols {
    dot1x {
      authenticator {
        authentication-profile-name profile1
        static [00:04:0f:fd:ac:fe 00:04:ae:cd:23:5f];
        interface {
          all {
            supplicant multiple;
          }
        }
      }
    }
  }
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying Static MAC Bypass of Authentication on page 2285

Verifying Static MAC Bypass of Authentication

Purpose Verify that the MAC address for both printers is configured and associated with the correct interfaces.

Action Use the operational mode command:

```
user@switch> show dot1x static-mac-address

MAC address           VLAN-Assignment          Interface
00:04:0f:fd:ac:fe     default                  ge-0/0/19.0
00:04:ae:cd:23:5f     default                  ge-0/0/20.0
```

Meaning The output field **MAC address** shows the MAC addresses of the two printers.

The output field **Interface** shows that the MAC address **00:04:0f:fd:ac:fe** can connect to the LAN through interface **ge-0/0/19.0** and that the MAC address **00:04:ae:cd:23:5f** can connect to the LAN through interface **ge-0/0/20.0**.

Related Documentation

- Configuring 802.1X Authentication (J-Web Procedure) on page 2332
- Configuring Static MAC Bypass of Authentication (CLI Procedure) on page 2334
- Configuring 802.1X Interface Settings (CLI Procedure) on page 2331
- Understanding MAC RADIUS Authentication on J-EX Series Switches

Example: Configuring MAC RADIUS Authentication on a J-EX Series Switch

To permit hosts that are not 802.1X-enabled to access the LAN, you can configure MAC RADIUS authentication on the switch interfaces to which the non-802.1X-enabled hosts are connected. When MAC RADIUS authentication is configured, the switch will attempt to authenticate the host with the RADIUS server using the host's MAC address.

This example describes how to configure MAC RADIUS authentication for two non-802.1X-enabled hosts:

Requirements

This example uses the following hardware and software components:

- One J-EX Series switch acting as an authenticator port access entity (PAE). The ports on the authenticator PAE form a control gate that blocks all traffic to and from supplicants until they are authenticated.
- One RADIUS authentication server. The authentication server acts as the backend database and contains credential information for hosts (supplicants) that have permission to connect to the network.

Before you configure MAC RADIUS authentication, be sure you have:

- Configured basic access between the J-EX Series switch and the RADIUS server. See "Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch" on page 2267.
- Performed basic bridging and VLAN configuration on the switch. See "Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch" on page 1063.
- Performed basic 802.1X configuration. See "Configuring 802.1X Interface Settings (CLI Procedure)" on page 2331.

Overview and Topology

IEEE 802.1X Port-Based Network Access Control (PNAC) authenticates and permits devices access to a LAN if the devices can communicate with the switch using the 802.1X protocol (are 802.1X-enabled). To permit non-802.1X-enabled hosts to access the LAN, you can configure MAC RADIUS authentication on the interfaces to which the hosts are connected. When the MAC address of the non-802.1X-enabled host appears on the interface, the switch consults the RADIUS server to check whether it is a permitted MAC address. If the MAC address of the host is configured as permitted on the RADIUS server, the switch opens LAN access to the nonresponsive host.

You can configure both MAC RADIUS authentication and 802.1X authentication methods on a single interface configured for multiple supplicants. Additionally, if an interface is

only connected to a non-802.1X-enabled host, you can enable MAC RADIUS and not enable 802.1X authentication using the **mac-radius restrict** option, and thus avoid the delay that occurs while the switch determines that the device is does not respond to EAP messages.

Figure 53 on page 2287 shows the two printers connected to the switch.

Figure 53: Topology for MAC RADIUS Authentication Configuration

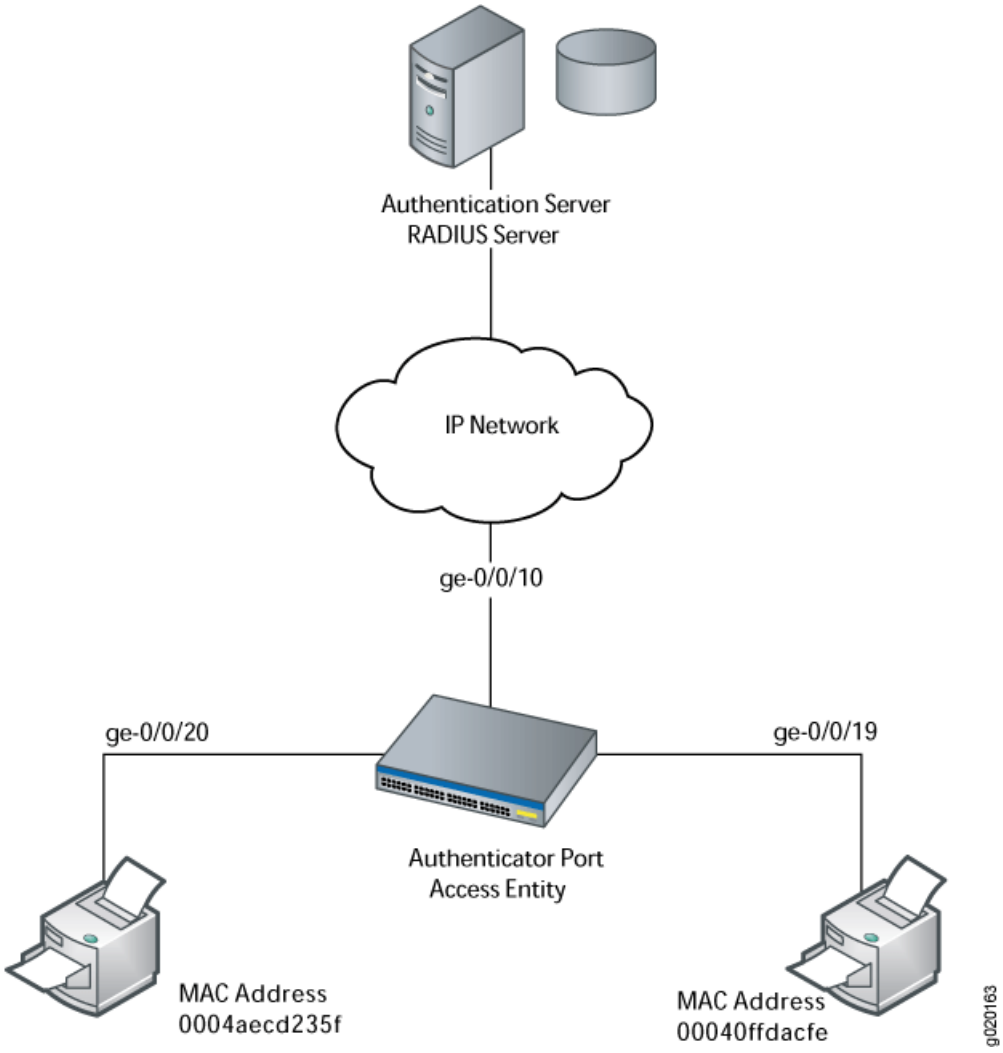

Table 292 on page 2287 shows the components in the example for MAC RADIUS authentication.

Table 292: Components of the MAC RADIUS Authentication Configuration Topology

Property	Settings
Switch hardware	J-EX4200 ports (**ge-0/0/0** through **ge-0/0/23**)
VLAN name	**default**

Table 292: Components of the MAC RADIUS Authentication Configuration Topology *(continued)*

Property	Settings
Connections to printers (no PoE required)	**ge-0/0/19**, MAC address 00040ffdacfe **ge-0/0/20**, MAC address 0004aecd235f
RADIUS server	Connected to the switch on interface **ge-0/0/10**

The printer with the MAC address 00040ffdacfe is connected to access interface **ge-0/0/19**. A second printer with the MAC address 0004aecd235f is connected to access interface **ge-0/0/20**. In this example, both interfaces are configured for MAC RADIUS authentication on the switch, and the MAC addresses (without colons) of both printers are configured on the RADIUS server. Interface **ge-0/0/20** is configured to eliminate the normal delay while the switch attempts 802.1X authentication; MAC RADIUS authentication is enabled and 802.1X authentication is disabled using the **mac-radius restrict** option.

Configuration

To configure MAC RADIUS authentication on the switch, perform these tasks:

CLI Quick Configuration

To quickly configure MAC RADIUS authentication, copy the following commands and paste them into the switch terminal window:

```
[edit]
set protocols dot1x authenticator interface ge-0/0/19 mac-radius
set protocols dot1x authenticator interface ge-0/0/20 mac-radius restrict
```

NOTE: You must also configure the two MAC addresses as usernames and passwords on the RADIUS server, as is done in Step 2 of the step-by-step procedure.

Step-by-Step Procedure

Configure MAC RADIUS authentication on the switch and on the RADIUS server:

1. On the switch, configure the interfaces to which the printers are attached for MAC RADIUS authentication, and configure interface **ge-0/0/20**, so that only MAC RADIUS authentication is used:

   ```
   [edit]
   user@switch# set protocols dot1x authenticator interface ge-0/0/19 mac-radius
   user@switch# set protocols dot1x authenticator interface ge-0/0/20 mac-radius restrict
   ```

2. On the RADIUS server, configure the MAC addresses **00040ffdacfe** and **0004aecd235f** as usernames and passwords:

   ```
   [root@freeradius]#
   edit /etc/raddb
   vi users
   00040ffdacfe Auth-type:=EAP, User-Password = "00040ffdacfe"
   0004aecd235f Auth-type:=EAP, User-Password = "0004aecd235f"
   ```

Results

Display the results of the configuration on the switch:

```
user@switch> show configuration
protocols {
  dot1x {
    authenticator {
      authentication-profile-name profile52;
      interface {
        ge-0/0/19.0 {
          mac-radius;
        }
        ge-0/0/20.0 {
          mac-radius {
            restrict;
          }
        }
      }
    }
  }
}
```

Verification

Verify that the supplicants are authenticated:

- Verifying That the Supplicants Are Authenticated on page 2289

Verifying That the Supplicants Are Authenticated

Purpose After supplicants are configured for MAC RADIUS authentication on the switch and on the RADIUS server, verify that they are authenticated and display the method of authentication:

Action Display information about 802.1X-configured interfaces **ge-0/0/19** and **ge-0/0/20**:

```
user@switch> show dot1x interface ge-0/0/19.0 detail
ge-0/0/19.0
  Role: Authenticator
  Administrative state: Auto
  Supplicant mode: Single
  Number of retries: 3
  Quiet period: 60 seconds
  Transmit period: 30 seconds
  Mac Radius: Enabled
  Mac Radius Restrict: Disabled
  Reauthentication: Enabled
  Configured Reauthentication interval: 3600 seconds
  Supplicant timeout: 30 seconds
  Server timeout: 30 seconds
  Maximum EAPOL requests: 2
  Guest VLAN member: <not configured>
  Number of connected supplicants: 1
    Supplicant: user101, 00:04:0f:fd:ac:fe
      Operational state: Authenticated
      Authentication method: Radius
      Authenticated VLAN: vo11
      Dynamic Filter: match source-dot1q-tag 10 action deny
```

```
      Session Reauth interval: 60 seconds
      Reauthentication due in 50 seconds

user@switch> show dot1x interface ge-0/0/20.0 detail
ge-0/0/20.0
  Role: Authenticator
  Administrative state: Auto
  Supplicant mode: Single
  Number of retries: 3
  Quiet period: 60 seconds
  Transmit period: 30 seconds
  Mac Radius: Enabled
  Mac Radius Restrict: Enabled
  Reauthentication: Enabled
  Configured Reauthentication interval: 3600 seconds
  Supplicant timeout: 30 seconds
  Server timeout: 30 seconds
  Maximum EAPOL requests: 2
  Guest VLAN member: <not configured>
  Number of connected supplicants: 1
    Supplicant: user102, 00:04:ae:cd:23:5f
      Operational state: Authenticated
      Authentcation method: Radius
      Authenticated VLAN: vo11
      Dynamic Filter: match source-dot1q-tag 10 action deny
      Session Reauth interval: 60 seconds
      Reauthentication due in 50 seconds
```

Meaning The sample output from the **show dot1x interface detail** command displays the MAC address of the connected host in the **Supplicant** field. On interface **ge-0/0/19**, the MAC address is **00:04:0f:fd:ac:fe**, which is the MAC address of the first printer configured for MAC RADIUS authentication. The **Authentication method** field displays the authentication method as **MAC Radius**. On interface **ge-0/0/20**, the MAC address is **00:04:ae:cd:23:5f**, which is the MAC address of the second printer configured for MAC RADIUS authentication. The **Authentication method** field displays the authentication method as **MAC Radius**.

Related Documentation

- Configuring MAC RADIUS Authentication (CLI Procedure) on page 2335
- Configuring 802.1X Interface Settings (CLI Procedure) on page 2331
- Configuring 802.1X Authentication (J-Web Procedure) on page 2332
- Understanding MAC RADIUS Authentication on J-EX Series Switches

Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on a J-EX Series Switch

802.1x Port-Based Network Access Control (PNAC) authentication on J-EX Series switches provides three types of authentication to meet the access needs of your enterprise LAN:

- Authenticate the first host (supplicant) on an authenticator port, and allow all others also connecting to have access.
- Authenticate only one supplicant on an authenticator port at one time.

- Authenticate multiple supplicants on an authenticator port. Multiple supplicant mode is used in VoIP configurations.

This example configures a J-EX4200 switch to use IEEE 802.1X to authenticate supplicants that use three different administrative modes:

Requirements

This example uses the following hardware and software components:

- One J-EX Series switch acting as an authenticator port access entity (PAE). The ports on the authenticator PAE form a control gate that blocks all traffic to and from supplicants until they are authenticated.
- One RADIUS authentication server that supports 802.1X. The authentication server acts as the backend database and contains credential information for hosts (supplicants) that have permission to connect to the network.

Before you configure the ports for 802.1X authentication, be sure you have:

- Installed your J-EX Series switch.
- Performed the initial switch configuration. See “Connecting and Configuring a J-EX Series Switch (J-Web Procedure)” on page 163.
- Performed basic bridging and VLAN configuration on the switch. See “Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch” on page 1063.
- Configured users on the authentication server.

Overview and Topology

As shown in Figure 54 on page 2292, the topology contains a J-EX4200 access switch connected to the authentication server on port **ge-0/0/10**. Interfaces **ge-0/0/8**, **ge-0/0/9**, and **ge-0/0/11** will be configured for three different administrative modes.

Figure 54: Topology for Configuring Supplicant Modes

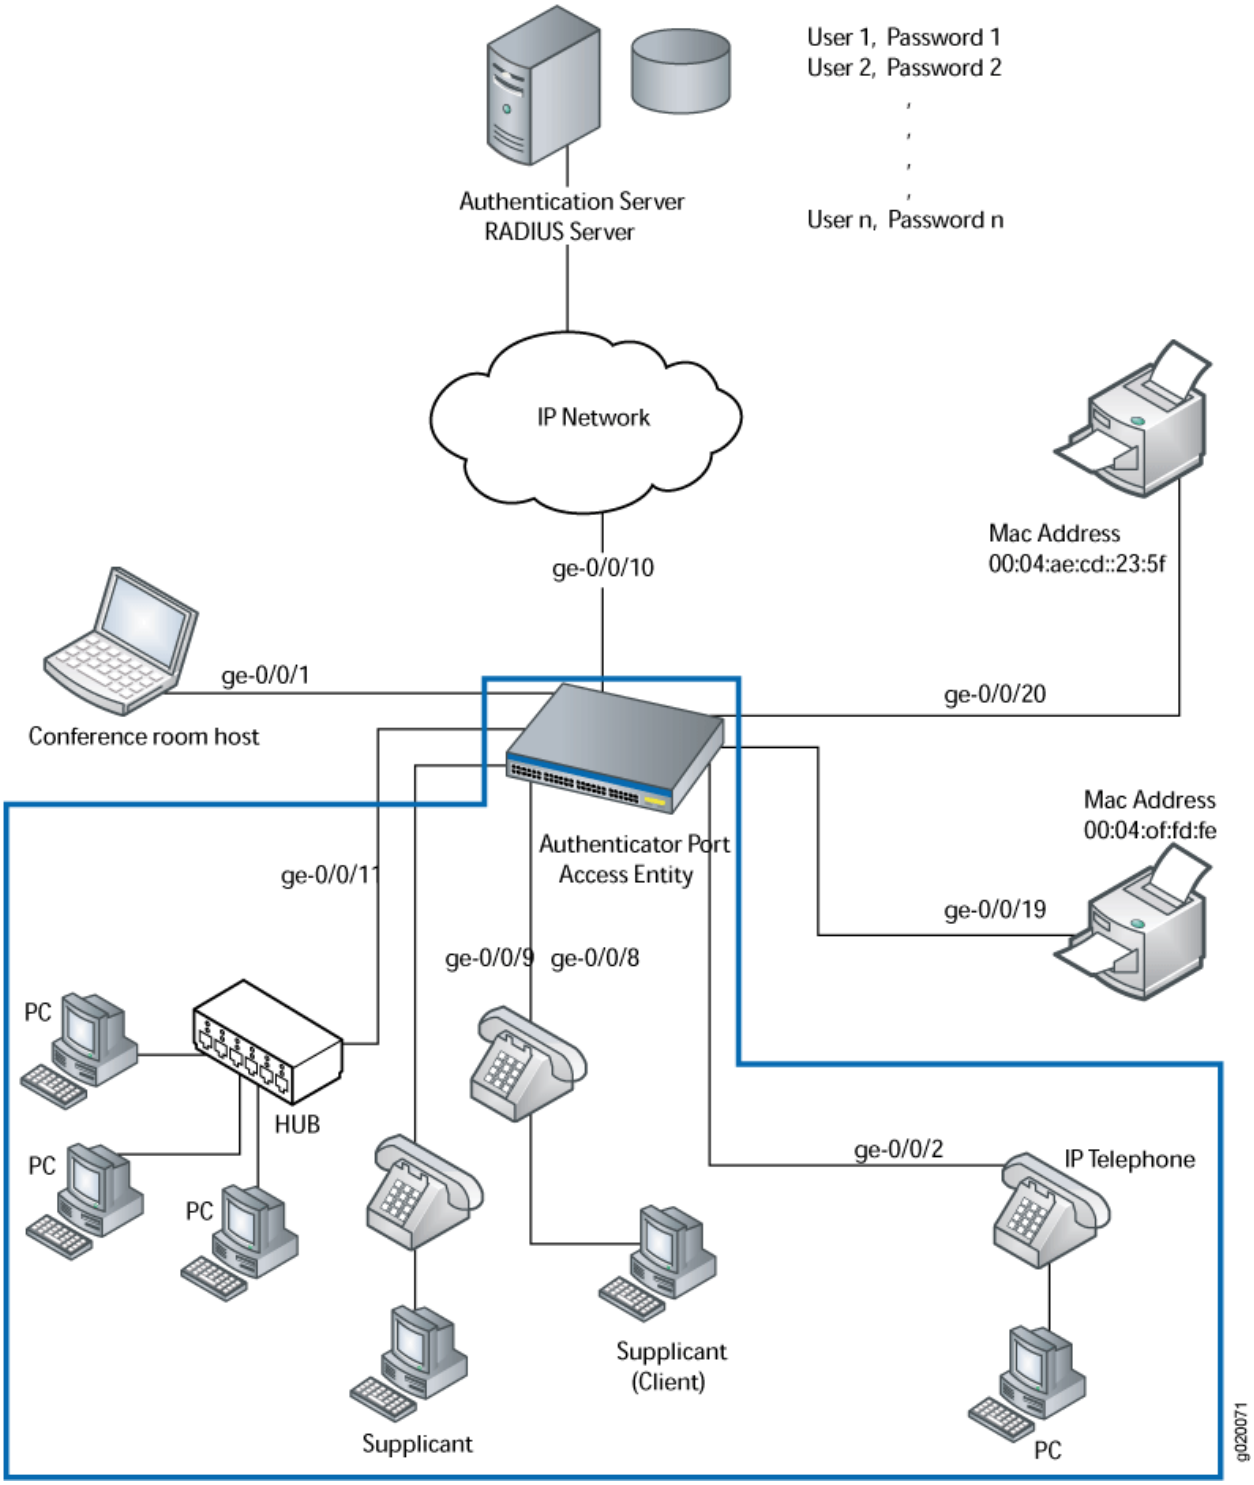

Table 293: Components of the Supplicant Mode Configuration Topology

Property	Settings
Switch hardware	J-EX4200 switch, 24 Gigabit Ethernet ports: 8 PoE ports (**ge-0/0/0** through **ge-0/0/7**) and 16 non-PoE ports (**ge-0/0/8** through **ge-0/0/23**)
Connections to Avaya phones—with integrated hub, to connect phone and desktop PC to a single port; (requires PoE)	**ge-0/0/8**, **ge-0/0/9**, and **ge-0/0/11**

To configure the administrative modes to support supplicants in different areas of the Enterprise network:

- Configure access port **ge-0/0/8** for single supplicant mode authentication.
- Configure access port **ge-0/0/9** for single secure supplicant mode authentication.
- Configure access port **ge-0/0/11** for multiple supplicant mode authentication.

Single supplicant mode authenticates only the first supplicant that connects to an authenticator port. All other supplicants connecting to the authenticator port after the first supplicant has connected successfully, whether they are 802.1X-enabled or not, are permitted free access to the port without further authentication. If the first authenticated supplicant logs out, all other supplicants are locked out until a supplicant authenticates.

Single-secure supplicant mode authenticates only one supplicant to connect to an authenticator port. No other supplicant can connect to the authenticator port until the first supplicant logs out.

Multiple supplicant mode authenticates multiple supplicants individually on one authenticator port. If you configure a maximum number of devices that can be connected to a port through port security, the lesser of the configured values is used to determine the maximum number of supplicants allowed per port.

Configuration of 802.1X to Support Multiple Supplicant Modes

To configure 802.1X authentication to support multiple supplicants, perform these tasks:

CLI Quick Configuration

To quickly configure the ports with different 802.1X authentication modes, copy the following commands and paste them into the switch terminal window:

```
[edit]
set protocols dot1x authenticator interface ge-0/0/8 supplicant single
set protocols dot1x authenticator interface ge-0/0/9 supplicant single-secure
set protocols dot1x authenticator interface ge-0/0/11 supplicant multiple
```

Step-by-Step Procedure

Configure the administrative mode on the interfaces:

1. Configure the supplicant mode as single on interface **ge-0/0/8**:

   ```
   [edit protocols]
   user@switch# set dot1x authenticator interface ge-0/0/8 supplicant single
   ```

2. Configure the supplicant mode as single secure on interface **ge-0/0/9**:

```
[edit protocols]
user@switch# set dot1x authenticator interface ge-0/0/9 supplicant single-secure
```

3. Configure multiple supplicant mode on interface **ge-0/0/11**:

```
[edit protocols]
user@switch# set dot1x authenticator interface ge-0/0/11 supplicant multiple
```

Results Check the results of the configuration:

```
[edit]
user@access-switch> show configuration
protocols {
  dot1x {
    authenticator {
      interface {
        ge-0/0/8.0 {
          supplicant single;
        )
        ge-0/0/9.0 {
          supplicant single-secure;
        )
        ge-0/0/11.0 {
          supplicant multiple;
        )
      }
    }
  }
}
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying the 802.1X Configuration on page 2294

Verifying the 802.1X Configuration

Purpose Verify the 802.1X configuration on interfaces **ge-0/0/8**, **ge-0/0/9**, and **ge-0/0/5**.

Action Verify the 802.1X configuration with the operational mode command **show dot1x interface**:

```
user@switch> show dot1x interface ge-0/0/8.0 detail
ge-0/0/8.0
  Role: Authenticator
  Administrative state: Auto
  Supplicant mode: Single
  Number of retries: 3
  Quiet period: 60 seconds
  Transmit period: 30 seconds
  Mac Radius: Disabled
  Mac Radius Restrict: Disabled
  Reauthentication: Enabled
  Configured Reauthentication interval: 3600 seconds
  Supplicant timeout: 30 seconds
  Server timeout: 30 seconds
```

```
  Maximum EAPOL requests: 2
  Guest VLAN member: <not configured>

user@switch> show dot1x interface ge-0/0/9.0 detail
ge-0/0/9.0
  Role: Authenticator
  Administrative state: Auto
  Supplicant mode: Single-Secure
  Number of retries: 3
  Quiet period: 60 seconds
  Transmit period: 30 seconds
  Mac Radius: Disabled
  Mac Radius Restrict: Disabled
  Reauthentication: Enabled
  Configured Reauthentication interval: 3600 seconds
  Supplicant timeout: 30 seconds
  Server timeout: 30 seconds
  Maximum EAPOL requests: 2
  Guest VLAN member: <not configured>
  Number of connected supplicants: 0

user@switch> show dot1x interface ge-0/0/11.0 detail
ge-0/0/11.0
  Role: Authenticator
  Administrative state: Auto
  Supplicant mode: Multiple
  Number of retries: 3
  Quiet period: 60 seconds
  Transmit period: 30 seconds
  Mac Radius: Disabled
  Mac Radius Restrict: Disabled
  Reauthentication: Enabled
  Configured Reauthentication interval: 3600 seconds
  Supplicant timeout: 30 seconds
  Server timeout: 30 seconds
  Maximum EAPOL requests: 2
  Guest VLAN member: <not configured>
  Number of connected supplicants: 0
```

Meaning The **Supplicant mode** output field displays the configured administrative mode for each interface. Interface **ge-0/0/8.0** displays **Single** supplicant mode. Interface **ge-0/0/9.0** displays **Single Secure** supplicant mode. Interface **ge-0/0/11.0** displays **Multiple** supplicant mode.

Related Documentation

- Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch on page 2267
- Example: Setting Up 802.1X in Conference Rooms to Provide Internet Access to Corporate Visitors on a J-EX Series Switch on page 2276
- Example: Setting Up VoIP with 802.1X and LLDP-MED on a J-EX Series Switch on page 2302
- Configuring 802.1X RADIUS Accounting (CLI Procedure) on page 2339
- Filtering 802.1X Supplicants Using RADIUS Server Attributes on page 2340
- Understanding Authentication onJ-EX Series Switches on page 2248

Example: Applying a Firewall Filter to 802.1X-Authenticated Supplicants Using RADIUS Server Attributes on a J-EX Series Switch

You can use RADIUS server attributes and a port-based firewall filter to centrally apply terms to multiple supplicants connected to a J-EX Series switch in your enterprise. Terms are applied following a supplicant's successful authentication through 802.1X.

J-EX Series switches support port-based firewall filters. Port firewall filters are configured on a single J-EX Series switch, but in order for them to operate throughout an enterprise, they have to be configured on multiple switches. To reduce the need to configure the same port firewall filter on multiple switches, you can instead apply the filter centrally on the RADIUS server using RADIUS server attributes.

The following example uses FreeRADIUS to apply a port firewall filter on a RADIUS server. For specifics on configuring your server, consult the documentation that was included with your RADIUS server.

This example describes how to configure a port firewall filter with terms, create counters to count packets for the supplicants, apply the filter to user profiles on the RADIUS server, and display the counters to verify the configuration:

Requirements

This example uses the following hardware and software components:

- One J-EX Series switch acting as an authenticator port access entity (PAE). The ports on the authenticator PAE form a control gate that blocks all traffic to and from supplicants until they are authenticated.
- One RADIUS authentication server. The authentication server acts as the backend database and contains credential information for hosts (supplicants) that have permission to connect to the network.

Before you connect the server to the switch, be sure you have:

- Set up a connection between the switch and the RADIUS server. See "Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch" on page 2267.
- Configured 802.1X authentication on the switch, with the authentication mode for interface **ge-0/0/2** set to **multiple**. See "Configuring 802.1X Interface Settings (CLI Procedure)" on page 2331 and "Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on a J-EX Series Switch" on page 2290.

- Configured users on the RADIUS authentication server (in this example, the user profiles for Supplicant 1 and Supplicant 2 in the topology are modified on the RADIUS server).

Overview and Topology

When the 802.1X configuration on an interface is set to **multiple** supplicant mode, you can apply a single port firewall filter configured through the Junos OS CLI on the J-EX Series switch to any number of users (supplicants) on one interface by adding the filter centrally to the RADIUS server. Only a single filter can be applied to an interface; however, the filter can contain multiple terms for separate supplicants.

For more information about firewall filters, see “Firewall Filters for J-EX Series Switches Overview” on page 2721.

RADIUS server attributes are applied to supplicants after the supplicants are successfully authenticated using 802.1X. To authenticate the supplicants, the switch forwards a supplicant’s credentials to the RADIUS server. The RADIUS server matches the credentials forwarded by the switch against preconfigured information about the supplicant located in the supplicant’s user profile on the RADIUS server. If a match is made, the RADIUS server instructs the switch to open an interface to the supplicant. Traffic then flows from and to the supplicant on the LAN. Further instructions configured in the port firewall filter and added to the supplicant’s user profile using a RADIUS server attribute further define the access that the supplicant is granted. Filtering terms configured in the port firewall filter are applied to the supplicant after 802.1X authentication is complete.

Figure 55 on page 2298 shows the topology used for this example. The RADIUS server is connected to the J-EX4200 switch on access port **ge-0/0/10**. Two supplicants are accessing the LAN on interface **ge-0/0/2**. Supplicant 1 has a MAC address of **00:50:8b:6f:60:3a**. Supplicant 2 has a MAC address of **00:50:8b:6f:60:3b**.

Figure 55: Topology for Firewall Filter and RADIUS Server Attributes Configuration

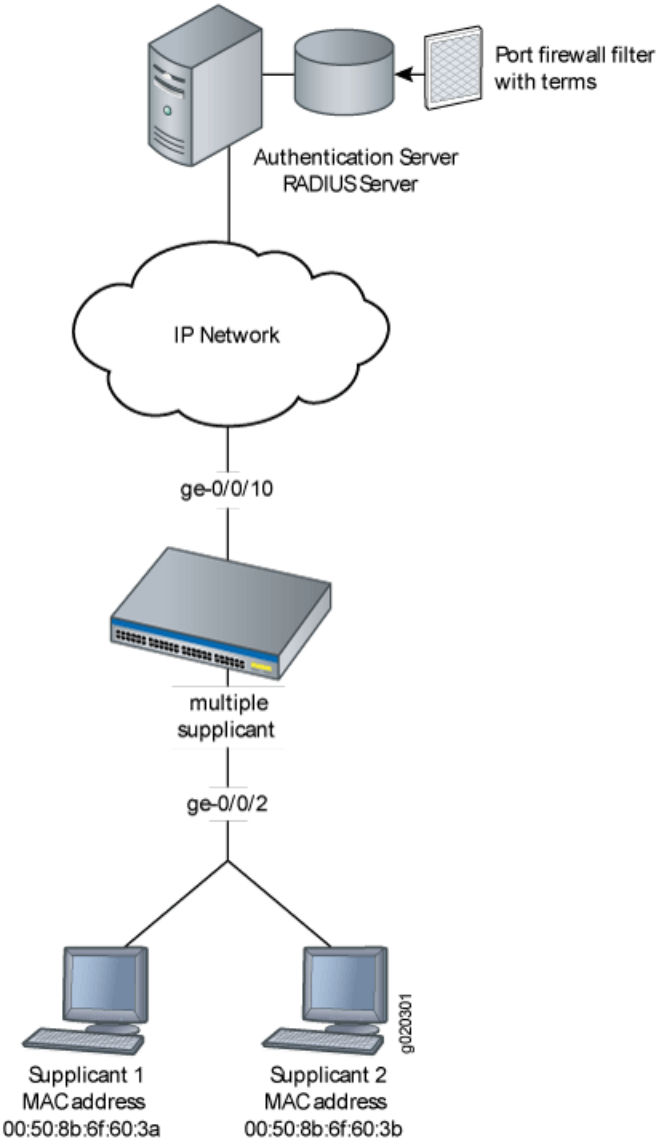

Table 294 on page 2298 describes the components in this topology.

Table 294: Components of the Firewall Filter and RADIUS Server Attributes Topology

Property	Settings
Switch hardware	J-EX4200 access switch, 24 Gigabit Ethernet ports, 8 PoE ports.
One RADIUS server	Backend database with an address of **10.0.0.100** connected to the switch at port **ge-0/0/10**.
802.1X supplicants connected to the switch on interface **ge-0/0/2**	• **Supplicant 1** has MAC address **00:50:8b:6f:60:3a**. • **Supplicant 2** has MAC address **00:50:8b:6f:60:3b**.

Table 294: Components of the Firewall Filter and RADIUS Server Attributes Topology *(continued)*

Property	Settings
Port firewall filter to be applied on the RADIUS server	**filter1**
Counters	**counter1** counts packets from Supplicant 1, and **counter2** counts packets from Supplicant 2.
User profiles on the RADIUS server	• Supplicant 1 has the user profile **supplicant1**. • Supplicant 2 has the user profile **supplicant2**.

In this example, you configure a port firewall filter named **filter1**. The filter contains terms that will be applied to the supplicants based on the MAC addresses of the supplicants. When you configure the filter, you also configure the counters called **counter1** and **counter2**. Packets from each supplicant will be counted, helping you verify that the configuration is working. Then, you check to see that the RADIUS server attribute is available on the RADIUS server and apply the filter to the user profiles of each supplicant on the RADIUS server. Finally, you verify the configuration by displaying output for the two counters.

NOTE: For more information about authentication, authorization, and accounting (AAA) services, see the *Junos OS System Basics Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.

Configuring the Port Firewall Filter and Counters

Configure a port firewall filter and counters:

CLI Quick Configuration

To quickly configure a port firewall filter with terms for Supplicant 1 and Supplicant 2 and create parallel counters for each supplicant, copy the following commands and paste them into the switch terminal window:

```
[edit]
set firewall family ethernet-switching filter filter1 term supplicant1 from source-mac-address
00:50:8b:6f:60:3a
set firewall family ethernet-switching filter filter1 term supplicant2 from source-mac-address
00:50:8b:6f:60:3b
set firewall family ethernet-switching filter filter1 term supplicant1 then count counter1
set firewall family ethernet-switching filter filter1 term supplicant2 then count counter2
```

Step-by-Step Procedure

To configure a port firewall filter and counters on the switch:

1. Configure a port firewall filter (here, **filter1**) with terms for each supplicant based upon the MAC address of each supplicant:

   ```
   [edit firewall family ethernet-switching]

   user@switch# set filter filter1 term supplicant1 from source-mac-address
   00:50:8b:6f:60:3a
   user@switch# set filter filter1 term supplicant2 from source-mac-address
   00:50:8b:6f:60:3b
   ```

2. Create two counters that will count packets for each supplicant:

```
[edit firewall family ethernet-switching]

user@switch# set filter filter1 term supplicant1 then count counter1
user@switch# set filter filter1 term supplicant2 then count counter2
user@switch# set filter filter1 term supplicant2 then count counter2
```

Results Display the results of the configuration:

```
user@switch> show configuration
  firewall {
    family ethernet-switching {
      filter filter1 {
        term supplicant1 {
          from {
            source-mac-address {
              00:50:8b:6f:60:3a;
            }
          }
          then count counter1;
          then policer p1;
        }
        term supplicant2 {
          from {
            source-mac-address {
              00:50:8b:6f:60:3b;
            }
          }
          then count counter2;
        }
      }
    }
  }
  policer p1 {
    if-exceeding {
      bandwidth-limit 1m;
        burst-size-limit 1k;
      }
      then discard;
    }
```

Applying the Port Firewall Filter to the Supplicant User Profiles on the RADIUS Server

Verify that the RADIUS server attribute needed to apply a filter on the RADIUS server is on the server and apply the port firewall filter to each supplicant's user profile on the RADIUS server:

Step-by-Step Procedure To verify that the RADIUS server attribute **Filter-ID** is on the RADIUS server and to apply the filter to the user profiles:

1. Display the dictionary **dictionary.rfc2865** on the RADIUS server, and verify that the attribute **Filter-ID** is in the dictionary:

   ```
   [root@freeradius]# cd usr/share/freeradius/dictionary.rfc2865
   ```

2. Close the dictionary file.

3. Display the local user profiles of the supplicants to which you want to apply the filter (here, the user profiles are called **supplicant1** and **supplicant2**):

```
[root@freeradius]# cat /usr/local/etc/raddb/users
```

The output shows:

```
supplicant1 Auth-Type := EAP, User-Password == "supplicant1"
        Tunnel-Type = VLAN,
        Tunnel-Medium-Type = IEEE-802,
        Tunnel-Private-Group-Id = "1005"

supplicant2 Auth-Type := EAP, User-Password == "supplicant2"
        Tunnel-Type = VLAN,
        Tunnel-Medium-Type = IEEE-802,
        Tunnel-Private-Group-Id = "1005"
```

4. Apply the filter to both user profiles by adding the line **Filter-Id = "filter1"** to each profile, and then close the file:

```
[root@freeradius]# cat /usr/local/etc/raddb/users
```

After you paste the line into the files, the files look like this:

```
supplicant1 Auth-Type := EAP, User-Password == "supplicant1"
        Tunnel-Type = VLAN,
        Tunnel-Medium-Type = IEEE-802,
        Tunnel-Private-Group-Id = "1005",
        Filter-Id = "filter1"

supplicant2 Auth-Type := EAP, User-Password == "supplicant2"
        Tunnel-Type = VLAN,
        Tunnel-Medium-Type = IEEE-802,
        Tunnel-Private-Group-Id = "1005",
        Filter-Id = "filter1"
```

Verification

Verify that the filter has been applied to the supplicants:

- Verifying That the Filter Has Been Applied to the Supplicants on page 2301

Verifying That the Filter Has Been Applied to the Supplicants

Purpose After supplicants are authenticated, verify that the filter configured on the switch and added to each supplicant's user profile on the RADIUS server has been applied:

Action Display information about firewall filter **filter1**:

```
user@switch> show firewall filter filter1
Filter: filter1
Counters:
Name                                                Bytes              Packets
counter1                                              128                    2
counter2                                               64                    1
```

Meaning The output of the command **show firewall filter filter1** displays **counter1** and **counter2**. Packets from Supplicant 1 are counted using **counter1**, and packets from Supplicant 2 are counted using **counter2**. The output from the command displays packets incrementing for both counters. The filter has been applied to both supplicants.

Related Documentation

- Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on a J-EX Series Switch on page 2290
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Configuring 802.1X RADIUS Accounting (CLI Procedure) on page 2339
- Understanding Authentication onJ-EX Series Switches on page 2248
- Understanding 802.1X and VSAs on J-EX Series Switches on page 2266

Example: Setting Up VoIP with 802.1X and LLDP-MED on a J-EX Series Switch

You can configure voice over IP (VoIP) on a J-EX Series switch to support IP telephones. The Link Layer Discovery Protocol–Media Endpoint Discovery (LLDP-MED) protocol forwards VoIP parameters from the switch to the phone. You also configure 802.1X authentication to allow the telephone access to the LAN. Authentication is done through a backend RADIUS server.

This example describes how to configure VoIP on a J-EX Series switch to support an Avaya IP phone, as well as the LLDP-MED protocol and 802.1X authentication:

- Requirements on page 2302
- Overview and Topology on page 2303
- Configuration on page 2305
- Verification on page 2307

Requirements

This example uses the following hardware and software components:

- One J-EX Series switch acting as an authenticator port access entity (PAE). The interfaces on the authenticator PAE form a control gate that blocks all traffic to and from supplicants until they are authenticated.
- An Avaya 9620 IP telephone that supports LLDP-MED and 802.1X

Before you configure VoIP, be sure you have:

- Installed your J-EX Series switch. See Installing and Connecting a J-EX4200 Switch.
- Performed the initial switch configuration. See “Connecting and Configuring a J-EX Series Switch (J-Web Procedure)” on page 163.

- Performed basic bridging and VLAN configuration on the switch. See “Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch” on page 1063.
- Configured the RADIUS server for 802.1X authentication and set up the access profile. See “Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch” on page 2267.
- (Optional) Configured interface **ge-0/0/2** for Power over Ethernet (PoE). The PoE configuration is not necessary if the VoIP supplicant is using a power adapter. For information about configuring PoE, see “Configuring PoE (CLI Procedure)” on page 3021.

NOTE: If the IP address isn't configured on the Avaya IP phone, the phone exchanges LLDP-MED information to get the VLAN ID for the voice VLAN. You must configure the **voip** statement on the interface to designate the interface as a VoIP interface and allow the switch to forward the VLAN name and VLAN ID for the voice VLAN to the IP telephone. The IP telephone then uses the voice VLAN (that is, it references the voice VLAN's ID) to send a DHCP discover request and exchange information with the DHCP server (voice gateway).

Overview and Topology

Instead of using a regular telephone, you connect an IP telephone directly to the switch. An IP phone has all the hardware and software needed to handle VoIP. You also can power an IP telephone by connecting it to one of the Power over Ethernet (PoE) interfaces on the switch.

In this example, the access interface **ge-0/0/2** on the J-EX4200 switch is connected to an Avaya 9620 IP telephone. Avaya phones have a built-in bridge that allows you to connect a desktop PC to the phone, so the desktop and phone in a single office require only one interface on the switch. The J-EX Series switch is connected to a RADIUS server on interface **ge-0/0/10** (see Figure 56 on page 2304).

Figure 56: VoIP Topology

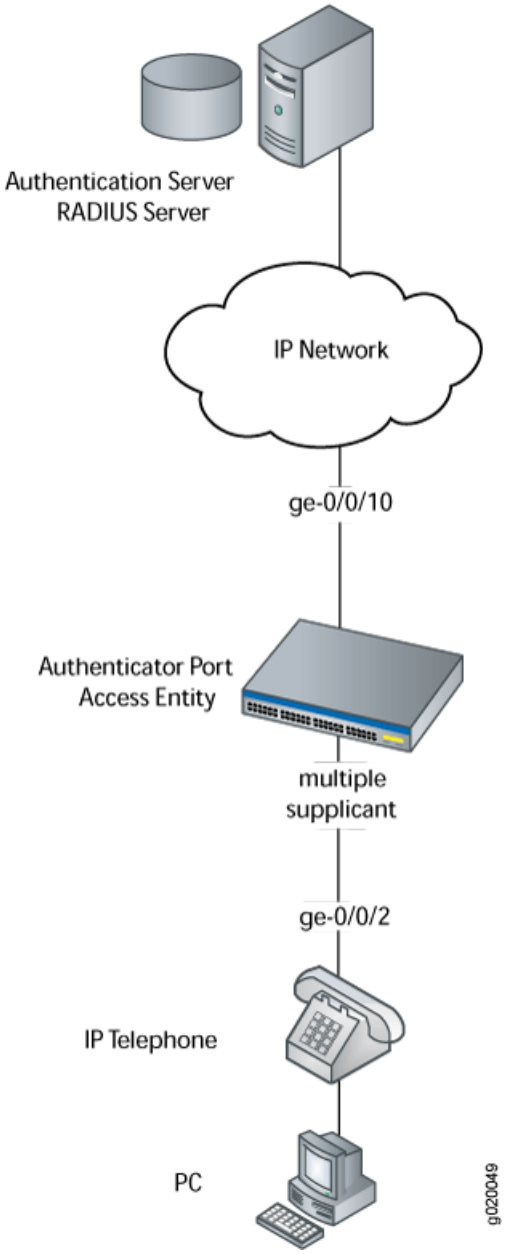

In this example, you configure VoIP parameters and specify the forwarding class **assured-forward** for voice traffic to provide the highest quality of service.

Table 295 on page 2304 describes the components used in this VoIP configuration example.

Table 295: Components of the VoIP Configuration Topology

Property	Settings
Switch hardware	J-EX4200 switch

Table 295: Components of the VoIP Configuration Topology *(continued)*

Property	Settings
VLAN names	**data-vlan** **voice-vlan**
Connection to Avaya phone—with integrated hub, to connect phone and desktop PC to a single interface (requires PoE)	**ge-0/0/2**
One RADIUS server	Provides backend database connected to the switch through interface **ge-0/0/10**.

As well as configuring a VoIP for interface **ge-0/0/2**, you configure:

- 802.1X authentication. Authentication is set to **multiple** supplicant to support more than one supplicant's access to the LAN through interface **ge-0/0/2**.
- LLDP-MED protocol information. The switch uses LLDP-MED to forward VoIP parameters to the phone. Using LLDP-MED ensures that voice traffic gets tagged and prioritized with the correct values at the source itself. For example, 802.1p class of service and 802.1Q tag information can be sent to the IP telephone.

NOTE: A PoE configuration is not necessary if an IP telephone is using a power adapter.

Configuration

To configure VoIP, LLDP-MED, and 802.1X authentication:

CLI Quick Configuration

To quickly configure VoIP, LLDP-MED, and 802.1X, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans data-vlan vlan-id 77
set vlans voice-vlan vlan-id 99
set vlans data-vlan interface ge-0/0/2.0
set interfaces ge-0/0/2 unit 0 family ethernet-switching vlan members data-vlan
set interfaces ge-0/0/2 unit 0 family ethernet-switching port-mode access
set ethernet-switching-options voip interface ge-0/0/2.0 vlan voice-vlan
set ethernet-switching-options voip interface ge-0/0/2.0 forwarding-class assured-forwarding
set protocols lldp-med interface ge-0/0/2.0
set protocols dot1x authenticator interface ge-0/0/2.0 supplicant multiple
```

Step-by-Step Procedure

To configure VoIP with LLDP-MED and 802.1X:

1. Configure the VLANs for voice and data:

```
[edit vlans]
user@switch# set data-vlan vlan-id 77
user@switch# set voice-vlan vlan-id 99
```

2. Associate the VLAN **data-vlan** with the interface:

```
[edit vlans]
user@switch# set data-vlan interface ge-0/0/2.0
```

3. Configure the interface as an access interface, configure support for Ethernet switching, and add the **data-vlan** VLAN:

```
[edit interfaces]
user@switch# set ge-0/0/2 unit 0 family ethernet-switching vlan members data-vlan
user@switch# set ge-0/0/2 unit 0 family ethernet-switching port-mode access
```

4. Configure VoIP on the interface and specify the **assured-forwarding** forwarding class to provide the most dependable class of service:

```
[edit ethernet-switching-options]
user@switch# set voip interface ge-0/0/2.0 vlan voice-vlan
user@switch# set voip interface ge-0/0/2.0 forwarding-class assured-forwarding
```

5. Configure LLDP-MED protocol support:

```
[edit protocols]
user@switch# set lldp-med interface ge-0/0/2.0
```

6. To authenticate an IP phone and a PC connected to the IP phone on the interface, configure 802.1X authentication support and specify **multiple** supplicant mode:

NOTE: If you do not want to authenticate any device, skip the 802.1X configuration on this interface.

```
[edit protocols]
user@switch# set dot1x authenticator interface ge-0/0/2.0 supplicant multiple
```

Results Display the results of the configuration:

```
[edit]
user@switch# show configuration
interfaces {
    ge-0/0/2 {
        unit 0 {
            family ethernet-switching {
                port-mode access;
                vlan {
                    members data-vlan;
                }
            }
        }
    }
}
protocols {
    lldp-med {
        interface ge-0/0/2.0;
    }
    dot1x {
        authenticator {
            interface {
                ge-0/0/2.0 {
                    supplicant multiple;
                }
            }
```

```
    }
  }
}
vlans {
  data-vlan {
    vlan-id 77;
    interface {
      ge-0/0/2.0;
    }
  }
  voice-vlan {
    vlan-id 99;
  }
}
ethernet-switching options {
  voip {
    interface ge-0/0/2.0 {
      vlan voice-vlan;
      forwarding-class assured-forwarding;
    }
  }
}
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying LLDP-MED Configuration on page 2307
- Verifying 802.1X Authentication for IP Phone and Desktop PC on page 2308
- Verifying the VLAN Association with the Interface on page 2309

Verifying LLDP-MED Configuration

Purpose Verify that LLDP-MED is enabled on the interface.

Action

```
user@switch> show lldp detail
LLDP                     : Enabled
Advertisement interval   : 30 Second(s)
Transmit delay           : 2 Second(s)
Hold timer               : 2 Second(s)
Config Trap Interval     : 300 Second(s)
Connection Hold timer    : 60 Second(s)

LLDP MED                 : Enabled
MED fast start count     : 3 Packet(s)


Interface      LLDP       LLDP-MED    Neighbor count
all            Enabled    -           0
ge-0/0/2.0     -          Enabled     0

Interface     VLAN-id     VLAN-name
ge-0/0/0.0    0           default
ge-0/0/1.0    0           employee-vlan
ge-0/0/2.0    0           data-vlan
ge-0/0/2.0    99          voice-vlan
ge-0/0/3.0    0           employee-vlan
```

```
ge-0/0/8.0    0          employee-vlan
ge-0/0/10.0   0          default
ge-0/0/11.0   20         employee-vlan
ge-0/0/23.0   0          default

LLDP basic TLVs supported:
Chassis identifier, Port identifier, Port description, System name, System
description, System capabilities, Management address.

LLDP 802 TLVs supported:
Power via MDI, Link aggregation, Maximum frame size, Port VLAN tag, Port
VLAN name.

LLDP MED TLVs supported:
LLDP MED capabilities, Network policy, Endpoint location, Extended power
Via MDI.
```

Meaning The **show lldp detail** output shows that both LLDP and LLDP-MED are configured on the **ge-0/0/2.0** interface. The end of the output shows the list of supported LLDP basic TLVs, 802.3 TLVs, and LLDP-MED TLVs that are supported.

Verifying 802.1X Authentication for IP Phone and Desktop PC

Purpose Display the 802.1X configuration to confirm that the VoIP interface has access to the LAN.

Action

```
user@switch> show dot1x interface ge/0/0/2.0 detail
ge-0/0/2.0
  Role: Authenticator
  Administrative state: Auto
  Supplicant mode: Multiple
  Number of retries: 3
  Quiet period: 60 seconds
  Transmit period: 30 seconds
  Mac Radius: Disabled
  Mac Radius Restrict: Disabled
  Reauthentication: Enabled
  Configured Reauthentication interval: 3600 seconds
  Supplicant timeout: 30 seconds
  Server timeout: 30 seconds
  Maximum EAPOL requests: 2
  Guest VLAN member: <not configured>
  Number of connected supplicants: 1
    Supplicant: user101, 00:04:0f:fd:ac:fe
      Operational state: Authenticated
      Authentication method: Radius
      Authenticated VLAN: vo11
      Dynamic Filter: match source-dot1q-tag 10 action deny
      Session Reauth interval: 60 seconds
      Reauthentication due in 50 seconds
```

Meaning The field **Role** shows that the **ge-0/0/2.0** interface is in the authenticator state. The **Supplicant** field shows that the interface is configured in multiple supplicant mode, permitting multiple supplicants to be authenticated on this interface. The MAC addresses of the supplicants currently connected are displayed at the bottom of the output.

Verifying the VLAN Association with the Interface

Purpose Display the interface state and VLAN membership.

Action

```
user@switch> show ethernet-switching interfaces
 Ethernet-switching table: 0 entries, 0 learned

user@switch> show ethernet-switching interfaces
Interface   State    VLAN members           Blocking
ge-0/0/0.0  down     default                unblocked
ge-0/0/1.0  down     employee-vlan          unblocked
ge-0/0/5.0  down     employee-vlan          unblocked
ge-0/0/3.0  down     employee-vlan          unblocked
ge-0/0/8.0  down     employee-vlan          unblocked
ge-0/0/10.0 down     default                unblocked
ge-0/0/11.0 down     employee-vlan          unblocked
ge-0/0/23.0 down     default                unblocked
ge-0/0/2.0  up       voice-vlan             unblocked
                     data-vlan              unblocked
```

Meaning The field **VLAN members** shows that the **ge-0/0/2.0** interface supports both the **data-vlan** VLAN and **voice-vlan** VLAN. The **State** field shows that the interface is up.

Related Documentation

- Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch on page 2267
- Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on a J-EX Series Switch on page 2290
- Defining CoS Forwarding Classes (CLI Procedure) on page 2918
- Defining CoS Forwarding Classes (J-Web Procedure) on page 2918
- Configuring LLDP-MED (CLI Procedure) on page 2346

Example: Configuring VoIP on a J-EX Series Switch Without Including 802.1X Authentication

You can configure voice over IP (VoIP) on a J-EX Series switch to support IP telephones.

To configure VoIP on a J-EX Series switch to support an IP phone that does not support 802.1X authentication, you must either add the MAC address of the phone to the static MAC bypass list or enable MAC RADIUS authentication on the switch.

This example describes how to configure VoIP on a J-EX Series switch without 802.1X authentication using static MAC bypass of authentication:

- Requirements on page 2310
- Overview on page 2310
- Configuration on page 2310
- Verification on page 2313

Requirements

This example uses the following hardware and software components:

- An IP telephone

Before you configure VoIP, be sure you have:

- Installed your J-EX Series switch. See Installing and Connecting a J-EX4200 Switch.
- Performed the initial switch configuration. See “Connecting and Configuring a J-EX Series Switch (J-Web Procedure)” on page 163.
- Performed basic bridging and VLAN configuration on the switch. See “Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch” on page 1063.
- Configured the RADIUS server for 802.1X authentication and set up the access profile. See “Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch” on page 2267.
- (Optional) Configured interface **ge-0/0/2** for Power over Ethernet (PoE). The PoE configuration is not necessary if the VoIP supplicant is using a power adapter. For information about configuring PoE, see “Configuring PoE (CLI Procedure)” on page 3021.

NOTE: If the IP address isn't configured on the Avaya IP phone, the phone exchanges LLDP-MED information to get the VLAN ID for the voice VLAN. You must configure the **voip** statement on the interface to designate the interface as a VoIP interface and allow the switch to forward the VLAN name and VLAN ID for the voice VLAN to the IP telephone. The IP telephone then uses the voice VLAN (that is, it references the voice VLAN’s ID) to send a DHCP discover request and exchange information with the DHCP server (voice gateway).

Overview

Instead of using a regular telephone, you connect an IP telephone directly to the switch. An IP phone has all the hardware and software needed to handle VoIP. You also can power an IP telephone by connecting it to one of the Power over Ethernet (PoE) interfaces on the switch.

In this example, the access interface **ge-0/0/2** on the J-EX4200 switch is connected to a non-802.1X IP phone.

To configure VoIP on a J-EX Series switch to support an IP phone that does not support 802.1X authentication, add the MAC address of the phone as a static entry in the authenticator database and set the supplicant mode to multiple.

Configuration

To configure VoIP without 802.1X authentication:

CLI Quick Configuration

To quickly configure VoIP, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans data-vlan vlan-id 77
set vlans voice-vlan vlan-id 99
set vlans data-vlan interface ge-0/0/2.0
set interfaces ge-0/0/2 unit 0 family ethernet-switching vlan members data-vlan
set interfaces ge-0/0/2 unit 0 family ethernet-switching port-mode access
set ethernet-switching-options voip interface ge-0/0/2.0 vlan voice-vlan
set ethernet-switching-options voip interface ge-0/0/2.0 forwarding-class assured-forwarding
set protocols lldp-med interface ge-0/0/2.0
set protocols dot1x authenticator authentication-profile-name auth-profile
set protocols dot1x authenticator static 00:04:f2:11:aa:a7
set protocols dot1x authenticator interface ge-0/0/2.0 supplicant multiple
```

Step-by-Step Procedure

To configure VoIP without 802.1X:

1. Configure the VLANs for voice and data:

   ```
   [edit vlans]
   user@switch# set data-vlan vlan-id 77
   user@switch# set voice-vlan vlan-id 99
   ```

2. Associate the VLAN **data-vlan** with the interface:

   ```
   [edit vlans]
   user@switch# set data-vlan interface ge-0/0/2.0
   ```

3. Configure the interface as an access interface, configure support for Ethernet switching, and add the **data-vlan** VLAN:

   ```
   [edit interfaces]
   user@switch# set ge-0/0/2 unit 0 family ethernet-switching vlan members data-vlan
   user@switch# set ge-0/0/2 unit 0 family ethernet-switching port-mode access
   ```

4. Configure VoIP on the interface and specify the **assured-forwarding** forwarding class to provide the most dependable class of service:

   ```
   [edit ethernet-switching-options]
   user@switch# set voip interface ge-0/0/2.0 vlan voice-vlan
   user@switch# set voip interface ge-0/0/2.0 forwarding-class assured-forwarding
   ```

5. Configure LLDP-MED protocol support:

   ```
   [edit protocols]
   user@switch# set lldp-med interface ge-0/0/2.0
   ```

6. Set the authentication profile (see “Configuring 802.1X Interface Settings (CLI Procedure)” on page 2331 and “Configuring 802.1X RADIUS Accounting (CLI Procedure)” on page 2339):

   ```
   [edit protocols]
   set dot1x authenticator authentication-profile-name auth-profile
   ```

7. Add the MAC address of the phone to the static MAC bypass list:

   ```
   [edit protocols]
   set dot1x authenticator static 00:04:f2:11:aa:a7
   ```

8. Set the supplicant mode to multiple:

```
[edit protocols]
set dot1x authenticator interface ge-0/0/2.0 supplicant multiple
```

Results Display the results of the configuration:

```
[edit]
user@switch# show configuration
interfaces {
  ge-0/0/2 {
    unit 0 {
      family ethernet-switching {
        port-mode access;
        vlan {
          members data-vlan;
        }
      }
    }
  }
}
protocols {
  lldp-med {
    interface ge-0/0/2.0;
  }
  dot1x {
    authenticator {
      authentication-profile-name auth-profile;
        static {
          00:04:f2:11:aa:a7;
        }
      }
      interface {
        ge-0/0/2.0 {
          supplicant multiple;
        }
    }
  }
}
vlans {
  data-vlan {
    vlan-id 77;
    interface {
      ge-0/0/2.0;
    }
  }
  voice-vlan {
    vlan-id 99;
  }
}
ethernet-switching options {
  voip {
    interface ge-0/0/2.0 {
      vlan voice-vlan;
      forwarding-class assured-forwarding;
    }
  }
```

```
}
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying LLDP-MED Configuration on page 2313
- Verifying Authentication for the Desktop PC on page 2314
- Verifying the VLAN Association with the Interface on page 2314

Verifying LLDP-MED Configuration

Purpose Verify that LLDP-MED is enabled on the interface.

Action

```
user@switch> show lldp detail
LLDP                     : Enabled
Advertisement interval   : 30 Second(s)
Transmit delay           : 2 Second(s)
Hold timer               : 2 Second(s)
Config Trap Interval     : 300 Second(s)
Connection Hold timer    : 60 Second(s)

LLDP MED                 : Enabled
MED fast start count     : 3 Packet(s)


Interface      LLDP       LLDP-MED    Neighbor count
all            Enabled    -           0
ge-0/0/2.0     -          Enabled     0

Interface     VLAN-id     VLAN-name
ge-0/0/0.0    0           default
ge-0/0/1.0    0           employee-vlan
ge-0/0/2.0    0           data-vlan
ge-0/0/2.0    99          voice-vlan
ge-0/0/3.0    0           employee-vlan
ge-0/0/8.0    0           employee-vlan
ge-0/0/10.0   0           default
ge-0/0/11.0   20          employee-vlan
ge-0/0/23.0   0           default

LLDP basic TLVs supported:
Chassis identifier, Port identifier, Port description, System name, System
description, System capabilities, Management address.

LLDP 802 TLVs supported:
Power via MDI, Link aggregation, Maximum frame size, Port VLAN tag, Port
VLAN name.

LLDP MED TLVs supported:
LLDP MED capabilities, Network policy, Endpoint location, Extended power
Via MDI.
```

Meaning The **show lldp detail** output shows that both LLDP and LLDP-MED are configured on the **ge-0/0/2.0** interface. The end of the output shows the list of supported LLDP basic TLVs, 802.3 TLVs, and LLDP-MED TLVs that are supported.

Verifying Authentication for the Desktop PC

Purpose Display the 802.1X configuration for the desktop PC connected to the VoIP interface through the IP phone.

Action

```
user@switch>  show dot1x interface ge/0/0/2.0 detail
ge-0/0/2.0
  Role: Authenticator
  Administrative state: Auto
  Supplicant mode: Multiple
  Number of retries: 3
  Quiet period: 60 seconds
  Transmit period: 30 seconds
  Mac Radius: Disabled
  Mac Radius Restrict: Disabled
  Reauthentication: Enabled
  Configured Reauthentication interval: 3600 seconds
  Supplicant timeout: 30 seconds
  Server timeout: 30 seconds
  Maximum EAPOL requests: 2
  Guest VLAN member: <not configured>
  Number of connected supplicants: 1
    Supplicant: user101, 00:04:0f:fd:ac:fe
      Operational state: Authenticated
      Authentication method: Radius
      Authenticated VLAN: vo11
      Dynamic Filter: match source-dot1q-tag 10 action deny
      Session Reauth interval: 60 seconds
      Reauthentication due in 50 seconds
```

Meaning The field **Role** shows that the **ge-0/0/2.0** interface is in the authenticator state. The **Supplicant** field shows that the interface is configured in multiple supplicant mode, permitting multiple supplicants to be authenticated on this interface. The MAC addresses of the supplicants currently connected are displayed at the bottom of the output.

Verifying the VLAN Association with the Interface

Purpose Display the interface state and VLAN membership.

Action

```
user@switch>  show ethernet-switching interfaces
 Ethernet-switching table: 0 entries, 0 learned

user@switch> show ethernet-switching interfaces
Interface   State    VLAN members           Blocking
ge-0/0/0.0  down     default                unblocked
ge-0/0/1.0  down     employee-vlan          unblocked
ge-0/0/5.0  down     employee-vlan          unblocked
ge-0/0/3.0  down     employee-vlan          unblocked
ge-0/0/8.0  down     employee-vlan          unblocked
ge-0/0/10.0 down     default                unblocked
ge-0/0/11.0 down     employee-vlan          unblocked
ge-0/0/23.0 down     default                unblocked
ge-0/0/2.0  up       voice-vlan             unblocked
                     data-vlan              unblocked
```

Meaning The field **VLAN members** shows that the **ge-0/0/2.0** interface supports both the **data-vlan** VLAN and **voice-vlan** VLAN. The **State** field shows that the interface is up.

Related Documentation

- Example: Setting Up VoIP with 802.1X and LLDP-MED on a J-EX Series Switch on page 2302
- Example: Configuring VoIP on a J-EX Series Switch Without Including LLDP-MED Support on page 2315
- Understanding 802.1X and VoIP on J-EX Series Switches on page 2263
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261

Example: Configuring VoIP on a J-EX Series Switch Without Including LLDP-MED Support

You can configure voice over IP (VoIP) on a J-EX Series switch to support IP telephones. The Link Layer Discovery Protocol–Media Endpoint Discovery (LLDP-MED) protocol is sometimes used with IP phones to forward VoIP parameters from the switch to the phone. Not all IP phones support LLDP-MED, however.

This example describes how to configure VoIP on a J-EX Series switch without LLDP-MED and without 802.1X:

- Requirements on page 2315
- Overview on page 2316
- Configuration on page 2316
- Verification on page 2317

Requirements

This example uses the following hardware and software components:

- One J-EX4200 switch acting as an authenticator port access entity (PAE). The interfaces on the authenticator PAE form a control gate that blocks all traffic to and from supplicants until they are authenticated.
- A IP phone that does not support LLDP-MED.

Before you configure VoIP, be sure you have:

- Performed basic bridging and VLAN configuration on the switch. See “Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch” on page 1063.
- Configured the IP phone as a member of the voice VLAN.
- (Optional) Configured interface **ge-0/0/2** for Power over Ethernet (PoE). The PoE configuration is not necessary if the VoIP supplicant is using a power adapter. See “Configuring PoE (CLI Procedure)” on page 3021.

Overview

Instead of using a regular telephone, you connect an IP telephone directly to the switch. An IP phone has all the hardware and software needed to handle VoIP. You also can power an IP telephone by connecting it to one of the Power over Ethernet (PoE) interfaces on the switch.

To configure VoIP on a J-EX Series switch to support an IP phone that does not support LLDP-MED, add the port to which you want to connect the IP phone as a member of the voice VLAN and configure the data VLAN as the native VLAN on the J-EX Series switch. This configuration ensures that the voice traffic and data traffic do not affect each other.

In this example, the interface ge-0/0/2 on the J-EX4200 switch is connected to a non-LLDP-MED IP phone.

NOTE: The implementation of a voice VLAN on an IP telephone is vendor-specific. Consult the documentation that came with your IP telephone for instructions on configuring a voice VLAN. For example, on an Avaya phone, you can ensure that the phone gets the correct VoIP VLAN ID even in the absence of LLDP-MED by enabling DHCP option 176.

Configuration

To configure VoIP without LLDP-MED or 802.1X authentication:

CLI Quick Configuration

To quickly configure VoIP, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans data-vlan vlan-id 77
set vlans voice-vlan vlan-id 99
set vlans data-vlan interface ge-0/0/2.0
set ethernet-switching-options voip interface ge-0/0/2.0 vlan voice-vlan
set ethernet-switching-options voip interface ge-0/0/2.0 forwarding-class assured-forwarding
set interfaces ge-0/0/2 unit 0 family ethernet-switching vlan members voice-vlan
set interfaces ge-0/0/2 unit 0 family ethernet-switching native-vlan-id data-vlan
```

Step-by-Step Procedure

Configure VoIP:

1. Configure the VLANs for data and voice:

   ```
   [edit vlans]
   user@switch# set data-vlan vlan-id 77
   user@switch# set voice-vlan vlan-id 99
   ```

2. Configure the VLAN **data-vlan** on the interface:

   ```
   [edit vlans]
   user@switch# set data-vlan interface ge-0/0/2.0
   ```

3. Configure VoIP on the interface and specify the **assured-forwarding** forwarding class to provide the most dependable class of service:

   ```
   [edit ethernet-switching-options]
   user@switch# set voip interface ge-0/0/2.0 vlan voice-vlan
   ```

```
user@switch# set voip interface ge-0/0/2.0 forwarding-class assured-forwarding
```

4. Add the interface as a member of the voice VLAN:

```
[edit interfaces]
set ge-0/0/2 unit 0 family ethernet-switching vlan members voice-vlan
```

5. Configure **data-vlan** as native to this trunk interface:

```
[edit interfaces]
user@switch# set ge-0/0/2 unit 0 family ethernet-switching native-vlan-id data-vlan
```

Results Display the results of the configuration:

```
[edit]
user@switch# show configuration
interfaces {
    ge-0/0/2 {
        unit 0 {
            family ethernet-switching {
                port-mode trunk;
                vlan {
                    members voice-vlan;
                }
                native-vlan-id data-vlan;
            }
        }
    }
}
vlans {
    data-vlan {
        vlan-id 77;
        interface {
            ge-0/0/2.0;
        }
    }
    voice-vlan {
        vlan-id 99;
    }
}
ethernet-switching options {
    voip {
        interface ge-0/0/2.0 {
            vlan voice-vlan;
            forwarding-class assured-forwarding;
        }
    }
}
```

Verification

To confirm that the configuration is working properly, perform the following task:

- Verifying the VLAN Association With the Interface on page 2318

Verifying the VLAN Association With the Interface

Purpose Display the interface state and VLAN membership.

Action

```
user@switch>  show ethernet-switching interfaces
 Ethernet-switching table: 0 entries, 0 learned

user@switch> show ethernet-switching interfaces
Interface   State    VLAN members           Blocking
ge-0/0/0.0  down     default                unblocked
ge-0/0/1.0  down     employee-vlan          unblocked
ge-0/0/5.0  down     employee-vlan          unblocked
ge-0/0/3.0  down     employee-vlan          unblocked
ge-0/0/8.0  down     employee-vlan          unblocked
ge-0/0/10.0 down     default                unblocked
ge-0/0/11.0 down     employee-vlan          unblocked
ge-0/0/23.0 down     default                unblocked
ge-0/0/2.0  up       voice-vlan             unblocked
                     data-vlan              unblocked
```

Meaning The field **VLAN members** shows that the **ge-0/0/2.0** interface supports both the **data-vlan** VLAN and **voice-vlan** VLAN. The **State** field shows that the interface is up.

Related Documentation

- Example: Setting Up VoIP with 802.1X and LLDP-MED on a J-EX Series Switch on page 2302
- Example: Configuring VoIP on a J-EX Series Switch Without Including 802.1X Authentication on page 2309
- Understanding 802.1X and VoIP on J-EX Series Switches on page 2263
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261

Example: Applying Firewall Filters to Multiple Supplicants on Interfaces Enabled for 802.1X or MAC RADIUS Authentication

On J-EX Series switches, firewall filters that you apply to interfaces enabled for 802.1X or MAC RADIUS authentication are dynamically combined with the per-user policies sent to the switch from the RADIUS server. The switch uses internal logic to dynamically combine the interface firewall filter with the user policies from the RADIUS server and create an individualized policy for each of the multiple users or nonresponsive hosts that are authenticated on the interface.

This example describes how dynamic firewall filters are created for multiple supplicants on an 802.1X-enabled interface (the same principles shown in this example apply to interfaces enabled for MAC RADIUS authentication):

- Requirements on page 2319
- Overview and Topology on page 2319
- Configuration on page 2321
- Verification on page 2322

Requirements

This example uses the following hardware and software components:

- One J-EX Series switch
- One RADIUS authentication server. The authentication server acts as the backend database and contains credential information for hosts (supplicants) that have permission to connect to the network.

Before you apply firewall filters to an interface for use with multiple supplicants, be sure you have:

- Set up a connection between the switch and the RADIUS server. See “Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch” on page 2267.
- Configured 802.1X authentication on the switch, with the authentication mode for interface **ge-0/0/2** set to **multiple**. See “Configuring 802.1X Interface Settings (CLI Procedure)” on page 2331 and “Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on a J-EX Series Switch” on page 2290.
- Configured users on the RADIUS authentication server.

Overview and Topology

When the 802.1X configuration on an interface is set to multiple supplicant mode, the system dynamically combines interface firewall filter with the user policies sent to the switch from the RADIUS server during authentication and creates separate terms for each user. Because there are separate terms for each user authenticated on the interface, you can, as shown in this example, use counters to view the activities of individual users that are authenticated on the same interface.

When a new user (or an nonresponsive host) is authenticated on an interface, the system adds a term to the firewall filter associated with the interface, and the term (policy) for each user is associated with the MAC address of the user. The term for each user is based on the user-specific filters set on the RADIUS server and the filters configured on the interface. For example, as shown in Figure 57 on page 2320, when User1 is authenticated by the J-EX Series switch, the system creates the firewall filter **dynamic-filter-example**. When User2 is authenticated, another term is added to the firewall filter, and so on.

Figure 57: Conceptual Model: Dynamic Filter Updated for Each New User

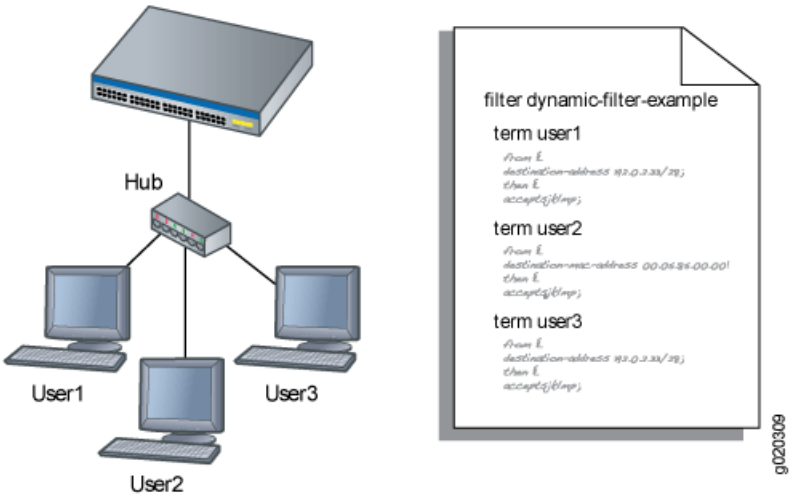

This is a conceptual model of the internal process—you cannot access or view the dynamic filter.

NOTE: If the firewall filter on the interface is modified after the user (or nonresponsive host) is authenticated, the modifications are not reflected in the dynamic filter unless the user is reauthenticated.

In this example, you configure a firewall filter to count the requests made by each endpoint authenticated on interface **ge-0/0/2** to the file server, which is located on subnet **192.0.2.16/28**. Figure 58 on page 2321 shows the network topology for this example.

Figure 58: Multiple Supplicants on an 802.1X-Enabled Interface Connecting to a File Server

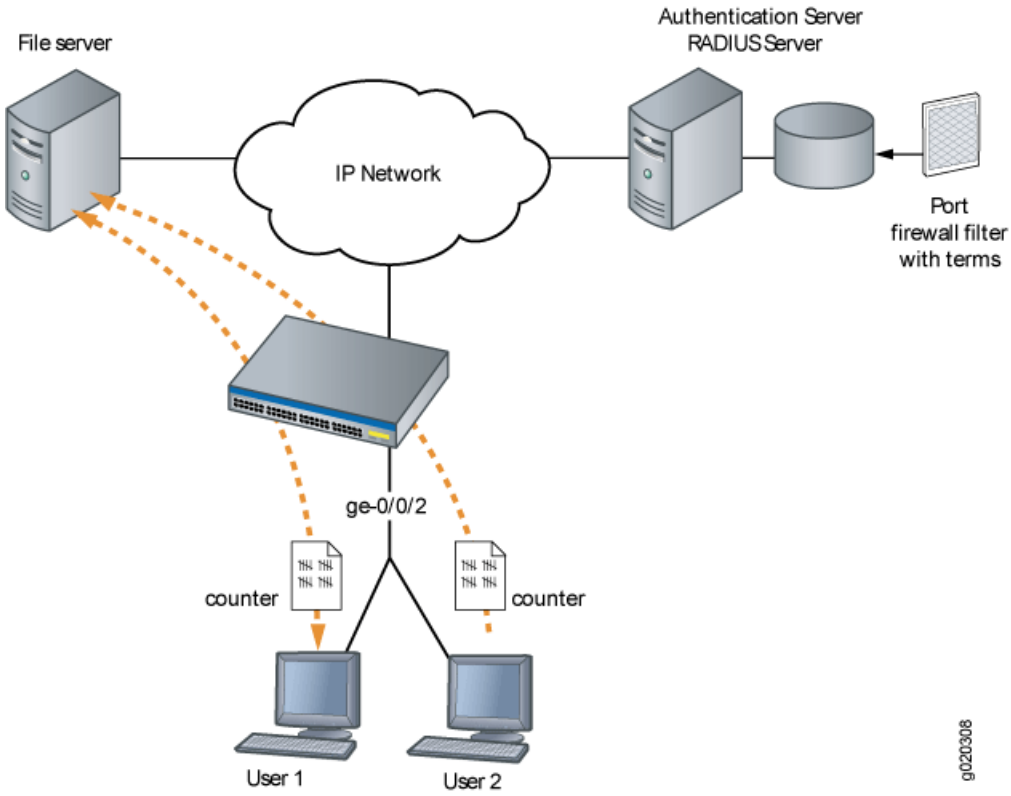

Configuration

To configure firewall filters for multiple supplicants on 802.1X-enabled interfaces:

- Configuring Firewall Filters on Interfaces with Multiple Supplicants on page 2321

Configuring Firewall Filters on Interfaces with Multiple Supplicants

CLI Quick Configuration

To quickly configure firewall filters on an interface enabled for multiple supplicants, copy the following commands and paste them into the switch terminal window:

```
[edit]
set protocols dot1x authenticator interface ge-0/0/2 supplicant multiple
set firewall family ethernet-switching filter filter1 term term1 from destination-address
192.0.2.16/28
set firewall family ethernet-switching filter filter1 term term1 then count counter1
```

Step-by-Step Procedure

To configure firewall filters on an interface enabled for multiple supplicants:

1. Configure interface **ge-0/0/2** for multiple supplicant mode authentication:

   ```
   [edit protocols dot1x]
   user@switch# set authenticator interface ge-0/0/2 supplicant multiple
   ```

2. Configure a firewall filter to count packets from each user. As each new user is authenticated on this multiple supplicant interface, this filter term will be included in the dynamically created term for the user:

```
[edit firewall family ethernet-switching]
user@switch# set filter filter1 term term1 from destination-address 192.0.2.16/28
user@switch# set filter filter1 term term1 then count counter1

user@switch# set filter filter1 term term2 then policer p1
```

Results Check the results of the configuration:

```
user@switch> show configuration

firewall {
    family ethernet-switching {
        filter filter1 {
            term term1 {
                from {
                    destination-address {
                        192.0.2.16/28;
                    }
                }
                then count counter1;
                term term2 {
                from {
                    destination-address {
                        192.0.2.16/28;
                    }
                }
                then policer p1;
            }
        }
    }
    policer p1 {
        if-exceeding {
            bandwidth-limit 1m;
                burst-size-limit 1k;
            }
            then discard;
        }
}
protocols {
    dot1x {
        authenticator
            interface ge-0/0/2 {
                supplicant multiple;
            }
        }
    }
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying Firewall Filters on Interfaces with Multiple Supplicants on page 2322

Verifying Firewall Filters on Interfaces with Multiple Supplicants

Purpose Verify that firewall filters are functioning on the interface with multiple supplicants.

Action

1. Check the results with one user authenticated on the interface. In this case, the user is authenticated on **ge-0/0/2**:

```
user@switch> show dot1x firewall

Filter: dot1x_ge-0/0/2
Counters
counter1_dot1x_ge-0/0/2_user1 100
```

2. When a second user, User2, is authenticated on the same interface, **ge-0/0/2**, you can verify that the filter includes the results for both of the users authenticated on the interface:

```
user@switch>  show dot1x firewall

Filter: dot1x-filter-ge-0/0/0
Counters
counter1_dot1x_ge-0/0/2_user1 100
counter1_dot1x_ge-0/0/2_user2 400
```

Meaning The results displayed by the **show dot1x firewall** output reflect the dynamic filter created with the authentication of each new user. User1 accessed the file server located at the specified destination address 100 times, while User2 accessed the same file server 400 times.

Related Documentation

- Example: Applying a Firewall Filter to 802.1X-Authenticated Supplicants Using RADIUS Server Attributes on a J-EX Series Switch on page 2296
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Filtering 802.1X Supplicants Using RADIUS Server Attributes on page 2340

Example: Setting Up Captive Portal Authentication on a J-EX Series Switch

You can set up captive portal authentication (hereafter referred to as captive portal) on a switch to redirect Web browser requests to a login page that requires the user to input a username and password. Upon successful authentication, the user is allowed to continue with the original page request and subsequent access to the network.

This example describes how to set up captive portal on a J-EX Series switch:

Requirements

This example uses the following hardware and software components:

- A J-EX4200 Series switch

Before you begin, be sure you have:

- Performed basic bridging and VLAN configuration on the switch. See “Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch” on page 1063.
- Generated an SSL certificate and installed it on the switch. See “Generating SSL Certificates to Be Used for Secure Web Access” on page 398.
- Configured basic access between the J-EX Series switch and the RADIUS server. See “Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch” on page 2267.
- Designed your captive portal login page. See “Designing and Modifying a Captive Portal Authentication Login Page on a J-EX Series Switch” on page 2351.

Overview and Topology

This example shows the configuration required on the switch to enable captive portal on an interface. To permit a printer connected to the captive portal interface to access the LAN, add its MAC address to the authentication whitelist. The MAC addresses on this list are permitted access on the interface without captive portal authentication.

The topology for this example consists of one J-EX Series switch connected to a RADIUS authentication server. One interface on the switch is configured for captive portal. In this example, the interface is configured in single supplicant mode.

Configuration

To configure captive portal on your switch:

CLI Quick Configuration

To quickly configure captive portal on the switch after completing the tasks in the Requirements section, copy the following commands and paste them into the switch terminal window:

```
[edit]
set system services web-management https local-certificate my-signed-cert
set services captive-portal secure-authentication https
set services captive-portal interface ge-0/0/10.0
set ethernet-switching-options authentication-whitelist 00:10:12:e0:28:22
set custom-options post-authentication-url http://www.my-home-page.com
```

Step-by-Step Procedure

To configure captive portal on the switch:

1. To create a secure channel for Web access to the switch, configure captive portal for HTTPS:

   a. Associate the security certificate with the Web server and enable HTTPS on the switch:

   ```
   [edit]
   user@switch# set system services web-management https local-certificate
   my-signed-cert
   ```

NOTE: You can enable HTTP instead of HTTPS, but we recommend HTTPS for security purposes.

b. Configure captive portal to use HTTPS:

```
[edit]
user@switch# set services captive-portal secure-authentication https
```

2. Enable an interface for captive portal:

```
[edit]
user@switch# set services captive-portal interface ge-0/0/10
```

3. (Optional) Allow specific clients to bypass captive portal authentication:

```
[edit]
user@switch# set ethernet-switching-options authentication-whitelist 00:10:12:e0:28:22
```

NOTE: Optionally, you can use set ethernet-switching-options authentication-whitelist 00:10:12:e0:28:22 interface ge-0/0/10.0 to limit the scope to the interface.

If the MAC address has already been learned on the interface, you must clear it using the **clear captive-portal interface *interface-name***) before adding it to the whitelist. Otherwise the new entry for the MAC address will not be added to the ethernet switching table and the authentication bypass will not be allowed.

4. (Optional) To redirect clients to a specified page rather than the page they originally requested, configure the post-authentication URL:

```
[edit services captive-portal]
user@switch# set custom-options post-authentication-url
http://www.my-home-page.com
```

Results Display the results of the configuration:

```
[edit]
user@switch# show
system {
    services {
        web-management {
            https {
                local-certificate my-signed-cert;
            }
        }
    }
}
security {
    certificates {
        local {
            my-signed-cert {
```

```
            "-----BEGIN RSA PRIVATE KEY-----\nMIICXwIBAAKBgQDk8sUggnXdDUmr7T
              vLv63yJq/LRpDASfIDZlX3z9ZDe1Kfk5C9\nr/tkyvzv
              ...
              Pt5YmvWDoGo0mSjoE/liH0BqYdh9YGqv3T2IEUfflSTQQHEOShS0ogWDHF\
              nnyOb1O/vQtjk20X9NVQg JHBwidssY9eRp\n-----END CERTIFICATE-----\n";
              ## SECRET-DATA
          }
        }
      }
    }
    services {
      captive-portal {
        interface {
          ge-0/0/10.0;
        }
        secure-authentication https;
      }
    }
    ethernet-switching-options {
      authentication-whitelist {
        00:10:12:e0:28:22/48;
      }
    }
```

Verification

To confirm that captive portal authentication is configured and working properly, perform these tasks:

- Verifying That Captive Portal Is Enabled on the Interface on page 2326
- Verify That Captive Portal Is Working Correctly on page 2326

Verifying That Captive Portal Is Enabled on the Interface

Purpose Verify that captive portal is configured on interface **ge-0/0/10**.

Action Use the operational mode command **show captive-portal interface *interface-name* detail**:

```
user@switch> show captive-portal interface ge-0/0/10.0 detail
ge-0/0/10.0
  Supplicant mode: Single
  Number of retries: 3
  Quiet period: 60 seconds
  Configured CP session timeout: 3600 seconds
  Server timeout: 15 seconds
```

Meaning The output confirms that captive portal is configured on interface ge-0/0/10 with the default settings for number of retries, quiet period, CP session timeout, and server timeout.

Verify That Captive Portal Is Working Correctly

Purpose Verify that captive portal is working on the switch.

Action Connect a client to interface **ge-0/0/10**. From the client, open a Web browser and request a webpage. The captive portal login page that you designed should be displayed. After

you enter your login information and are authenticated against the RADIUS server, the Web browser should display either the page you requested or the post-authentication URL that you configured.

Troubleshooting

To troubleshoot captive portal, perform these tasks:

Troubleshooting Captive Portal

Problem The switch does not return the captive portal login page when a user connected to a captive portal interface on the switch requests a Web page.

Solution You can examine the ARP, DHCP, HTTPS, and DNS counters—if one or more of these counters are not incrementing, this provides an indication of where the problem lies. For example, if the client cannot get an IP address, you might check the switch interface to determine whether the DHCP counter is incrementing—if the counter increments, the DHCP packet was received by the switch.

```
user@switch> show captive-portal firewall ge-0/0/10.0
ge-0/0/10.0
  Filter name: dot1x_ge-0/0/10
Counters:
Name                              Bytes              Packets
dot1x_ge-0/0/10_CP_arp             7616                  119
dot1x_ge-0/0/10_CP_dhcp               0                    0
dot1x_ge-0/0/10_CP_http               0                    0
dot1x_ge-0/0/10_CP_https              0                    0
dot1x_ge-0/0/10_CP_t_dns              0                    0
dot1x_ge-0/0/10_CP_u_dns              0                    0
```

Related Documentation

- Configuring Captive Portal Authentication (CLI Procedure) on page 2350Configuring
- Designing a Captive Portal Authentication Login Page on a J-EX Series Switch on page 2351

CHAPTER 83

Configuring Access Control

Specifying RADIUS Server Connections on a J-EX Series Switch (CLI Procedure)

IEEE 802.1X and MAC RADIUS authentication both provide network edge security, protecting Ethernet LANs from unauthorized user access by blocking all traffic to and from devices at the interface until the supplicant's credentials or MAC address are presented and matched on the *authentication server* (a RADIUS server). When the supplicant is authenticated, the switch stops blocking access and opens the interface to the supplicant.

To use 802.1X or MAC RADIUS authentication, you must specify the connections on the switch for each RADIUS server to which you will connect.

To configure a RADIUS server on the switch:

1. Define the IP address of the RADIUS server, the RADIUS server authentication port number, and the secret password. You can define more than one RADIUS server. The secret password on the switch must match the secret password on the server:

```
[edit access]
user@switch# set radius-server 10.0.0.100 port 1812 secret abc
```

NOTE: Specifying the authentication port is optional, and port 1812 is the default. However, we recommend that you configure it in order to avoid confusion as some RADIUS servers might refer to an older default.

2. (Optional) Specify the IP address by which the switch is identified by the RADIUS server. If you do not specify this, the RADIUS server uses the address of the interface sending the RADIUS request. We recommend that you specify this IP address because if the request gets diverted on an alternate route to the RADIUS server, the interface relaying the request might not be an interface on the switch.

```
[edit access]
user@switch# set access radius-erver source-address 10.93.14.100
```

3. Configure the authentication order, making **radius** the first method of authentication:

```
[edit access]
user@switch# set profile profile1 authentication-order radius
```

4. Create a profile and specify the list of RADIUS servers to be associated with the profile. For example, you might choose to group your RADIUS servers geographically by city. This feature enables easy modification whenever you want to change to a different sent of authentication servers.

```
[edit access profile]
user@switch# set atlanta radius authentication-server 10.0.0.100 10.2.14.200
```

5. Specify the group of servers to be used for 802.1X or MAC RADIUS authentication by identifying the profile name:

```
[edit access profile]
user@switch# set protocols dot1x authenticator authentication-profile-name denver
```

6. Configure the IP address of the J-EX Series switch in the list of clients on the RADIUS server. For specifics on configuring the RADIUS server, consult the documentation for your server.

Related Documentation

- Configuring 802.1X Interface Settings (CLI Procedure) on page 2331
- Configuring 802.1X Authentication (J-Web Procedure) on page 2332
- Configuring MAC RADIUS Authentication (CLI Procedure) on page 2335
- Configuring 802.1X RADIUS Accounting (CLI Procedure) on page 2339

Configuring 802.1X Interface Settings (CLI Procedure)

IEEE 802.1X authentication provides network edge security, protecting Ethernet LANs from unauthorized user access by blocking all traffic to and from a supplicant (client) at the interface until the supplicant's credentials are presented and matched on the *authentication server* (a RADIUS server). When the supplicant is authenticated, the switch stops blocking access and opens the interface to the supplicant.

NOTE: You can also specify an 802.1X exclusion list to specify supplicants can that can bypass authentication and be automatically connected to the LAN. See “Configuring Static MAC Bypass of Authentication (CLI Procedure)” on page 2334.

Before you begin, specify the RADIUS server or servers to be used as the authentication server. See “Specifying RADIUS Server Connections on a J-EX Series Switch (CLI Procedure)” on page 2330.

To configure 802.1X on an interface:

1. Configure the supplicant mode as **single** (authenticates the first supplicant), **single-secure** (authenticates only one supplicant), or **multiple** (authenticates multiple supplicants):

```
[edit protocols dot1x]
user@switch# set authenticator interface ge-0/0/5 supplicant multiple
```

2. Enable reauthentication and specify the reauthentication interval:

```
[edit protocols dot1x]
user@switch# set authenticator interface ge-0/0/5/0 reauthentication interval 5
```

3. Configure the interface timeout value for the response from the supplicant:

```
[edit protocols dot1x]
user@switch# set authenticator interface ge-0/0/5 supplicant-timeout 5
```

4. Configure the timeout for the interface before it resends an authentication request to the RADIUS server:

```
[edit protocols dot1x]
user@switch# set authenticator interface ge-0/0/5 server-timeout 5
```

5. Configure how long, in seconds, the interface waits before retransmitting the initial EAPOL PDUs to the supplicant:

```
[edit protocols dot1x]
user@switch# set authenticator interface ge-0/0/5 transmit-period 60
```

6. Configure the maximum number of times an EAPOL request packet is retransmitted to the supplicant before the authentication session times out:

```
[edit protocols dot1x]
user@switch# set authenticator interface ge-0/0/5 maximum-requests 5
```

Related Documentation

- Configuring 802.1X Authentication (J-Web Procedure) on page 2332
- Example: Setting Up VoIP with 802.1X and LLDP-MED on a J-EX Series Switch on page 2302
- Monitoring 802.1X Authentication on page 2355
- Verifying 802.1X Authentication on page 2356
- Configuring LLDP (CLI Procedure) on page 2344
- Understanding Authentication onJ-EX Series Switches on page 2248

Configuring 802.1X Authentication (J-Web Procedure)

To configure 802.1X settings on a J-EX Series switch using the J-Web interface:

1. Select **Configure** > **Security** > **802.1X**.

   The 802.1X screen displays a list of interfaces, whether 802.1X security has been enabled, and the assigned port role.

   When you select an interface, the **Details of 802.1x configuration on port** section displays 802.1X details for that interface.

   NOTE: After you make changes to the configuration in this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options** > **Commit**. See "Using the Commit Options to Commit Configuration Changes (J-Web Procedure)" on page 334 for details about all commit options.

2. Click one:

   - RADIUS Servers—Specifies the RADIUS server to be used for authentication. Select the check box to specify a server. Click **Add** or **Edit** to add or modify the RADIUS server settings. Enter information as specified in Table 296 on page 2333.
   - Exclusion List—Excludes hosts from the 802.1X authentication list by specifying the MAC address. Click **Add** or **Edit** in the Exclusion list screen to include or modify the MAC addresses. Enter information as specified in Table 297 on page 2333.
   - Edit—Specifies 802.1X settings for the selected interface

- Apply 802.1X Profile—Applies an 802.1X profile based on the port role. If a message appears asking whether you want to configure a RADIUS server, click **Yes**.
- 802.1X Configuration—Configures custom 802.1X settings for the selected interface. If a message appears asking if you want to configure a RADIUS server, click **Yes**. Enter information as specified in Table 296 on page 2333. To configure 802.1X settings, enter information as specified in Table 298 on page 2333.

- Delete—Deletes 802.1X authentication configuration on the selected interface.

Table 296: RADIUS Server Settings

Field	Function	Your Action
IP Address	Specifies the IP address of the server.	Enter the IP address in dotted decimal notation.
Password	Specifies the login password.	Enter the password.
Confirm Password	Verifies the login password for the server.	Reenter the password.
Server Port Number	Specifies the port with which the server is associated.	Type the port number.
Source Address	Specifies the source address of the switch using which the switch can communicate with the server.	Type the IP address in dotted decimal notation.
Retry Attempts	Specifies the number of login retries allowed after a login failure.	Type the number.
Timeout	Specifies the time interval to wait before the connection to the server is closed.	Type the interval in seconds.

Table 297: 802.1X Exclusion List

Field	Function	Your Action
MAC Address	Specifies the MAC address to be excluded from 802.1X authentication.	Enter the MAC address.
Exclude if connected through the port	Specifies that the host can bypass authentication if it is connected through a particular interface.	Select to enable the option. Select the port through which the host is connected.
Move the host to the VLAN	Specifies moving the host to a specific VLAN once the host is authenticated.	Select to enable the option. Select the VLAN from the list.

Table 298: 802.1X Port Settings

Field	Function	Your Action
Supplicant Mode		

Table 298: 802.1X Port Settings *(continued)*

Field	Function	Your Action
Supplicant Mode	Specifies the mode to be adopted for supplicants: • Single—allows only one host for authentication. • Multiple—allows multiple hosts for authentication. Each host is checked before being admitted to the network. • Single authentication for multiple hosts—Allows multiple hosts but only the first is authenticated.	Select a mode.
Authentication		
Enable re-authentication	Specifies enabling reauthentication on the selected interface.	1. Select to enable reauthentication. 2. Enter the timeout for reauthentication in seconds.
Action on authentication failure	Specifies the action to be taken in case the host does not respond, leading to an authentication failure.	Select one: • Move to the Guest VLAN—Select the VLAN to move the interface to. • Deny—The host is not permitted access.
Timeouts	Specifies timeout values for each action.	Enter the value in seconds for: • Port waiting time after an authentication failure • EAPOL retransmitting interval • Max. EAPOL requests • Maximum number of retries • Port timeout value for the response from the supplicant • Port timeout value for the response from the RADIUS server

Related Documentation

- Configuring 802.1X Interface Settings (CLI Procedure) on page 2331
- Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on a J-EX Series Switch on page 2290
- Understanding Authentication onJ-EX Series Switches on page 2248

Configuring Static MAC Bypass of Authentication (CLI Procedure)

You can configure a static MAC bypass list (sometimes called the exclusion list) on the switch to specify MAC addresses of devices allowed access to the LAN without 802.1X or MAC RADIUS authentication requests to the RADIUS server.

To configure the static MAC bypass list:

- Specify a MAC address to bypass authentication:

```
[edit protocols dot1x]
user@switch# set authenticator static 00:04:0f:fd:ac:fe
```

- Configure a supplicant to bypass authentication if connected through a particular interface:

```
[edit protocols dot1x]
user@switch# set authenticator static 00:04:0f:fd:ac:fe interface ge-0/0/5
```

- You can configure a supplicant to be moved to a specific VLAN after it is authenticated:

```
[edit protocols dot1x]
user@switch# set authenticator static 00:04:0f:fd:ac:fe interface ge-0/0/5
vlan-assignment default-vlan
```

Related Documentation

- Example: Configuring Static MAC Bypass of Authentication on a J-EX Series Switch on page 2281
- Configuring 802.1X Interface Settings (CLI Procedure) on page 2331
- Configuring 802.1X Authentication (J-Web Procedure) on page 2332

Configuring MAC RADIUS Authentication (CLI Procedure)

You can permit devices that are not 802.1X-enabled LAN access by configuring MAC RADIUS authentication on the J-EX Series switch interfaces to which the hosts are connected.

NOTE: You can also allow non-802.1X-enabled devices to access the LAN by configuring their MAC address for static MAC bypass of authentication.

You can configure MAC RADIUS authentication on an interface that also allows 802.1X authentication, or you can configure either authentication method alone.

If both MAC RADIUS and 802.1X authentication are enabled on the interface, the switch first sends the host three EAPOL requests to the host. If there is no response from the host, the switch sends the host's MAC address to the RADIUS server to check whether it is a permitted MAC address. If the MAC address is configured as permitted on the RADIUS server, the RADIUS server sends a message to the switch that the MAC address is a permitted address, and the switch opens LAN access to the nonresponsive host on the interface to which it is connected.

If MAC RADIUS authentication is configured on the interface but 802.1X authentication is not (by using the **mac-radius restrict** option), the switch attempts to authenticate the MAC address with the RADIUS server without delaying by attempting 802.1X authentication first.

Before you configure MAC RADIUS authentication, be sure you have:

- Configured basic access between the J-EX Series switch and the RADIUS server. See “Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch” on page 2267.

To configure MAC RADIUS authentication using the CLI:

- On the switch, configure the interfaces to which the nonresponsive hosts are attached for MAC RADIUS authentication, and add the **restrict** qualifier for interface **ge-0/0/20** to have it use only MAC RADIUS authentication:

```
[edit]
user@switch# set protocols dot1x authenticator interface ge-0/0/19 mac-radius
user@switch# set protocols dot1x authenticator interface ge-0/0/20 mac-radius restrict
```

- On a RADIUS authentication server, create user profiles for each nonresponsive host using the MAC address (without colons) of the nonresponsive host as the username and password (here, the MAC addresses are **00:04:0f:fd:ac:fe** and **00:04:ae:cd:23:5f**):

```
[root@freeradius]#
edit /etc/raddb
vi users
00040ffdacfe Auth-type:=Local, User-Password = "00040ffdacfe"
0004aecd235f Auth-type:=Local, User-Password = "0004aecd235f"
```

Related Documentation

- Example: Configuring MAC RADIUS Authentication on a J-EX Series Switch on page 2286
- Verifying 802.1X Authentication on page 2356
- Understanding MAC RADIUS Authentication on J-EX Series Switches

Configuring Server Fail Fallback (CLI Procedure)

Server fail fallback allows you to specify how end devices connected to the switch are supported if the RADIUS authentication server becomes unavailable or sends an Extensible Authentication Protocol Over LAN (EAPOL) access-reject message.

802.1X user authentication works by using an *authenticator port access entity* (the J-EX Series switch) to block all traffic to and from an end device at the interface until the end device's credentials are presented and matched on the *authentication server* (a RADIUS server). When the end device has been authenticated, the switch stops blocking and opens the interface to the end device.

When you set up 802.1X or MAC RADIUS authentication on the switch, you specify a primary authentication server and one or more backup authentication servers. If the primary authentication server cannot be reached by the switch and the secondary authentication servers are also unreachable, a RADIUS server timeout occurs. Because the authentication server grants or denies access to the end devices awaiting authentication, the switch does not receive access instructions for end devices attempting access to the LAN and normal authentication cannot be completed. Server fail fallback allows you to configure authentication alternatives that permit the switch to take appropriate actions toward end devices awaiting authentication or reauthentication.

To configure basic server fail fallback options using the CLI:

- Configure an interface to allow traffic to flow from a supplicant to the LAN if a RADIUS server timeout occurs (as if the supplicant had been successfully authenticated by a RADIUS server):

```
[edit protocols dot1x authenticator]
user@switch# set interface ge-0/0/1 server-fail permit
```

- Configure an interface to prevent traffic flow from an end device to the LAN (as if the end device had failed authentication and had been rejected by the RADIUS server):

```
[edit protocols dot1x authenticator]
user@switch# set interface ge-0/0/1 server-fail deny
```

- Configure an interface to move an end device to a specified VLAN if a RADIUS server timeout occurs (in this case, the VLAN name is **vlan1**):

```
[edit protocols dot1x authenticator]
user@switch# set interface ge-0/0/1 server-fail vlan-name vlan1
```

- Configure an interface to recognize already connected end devices as reauthenticated if there is a RADIUS timeout during reauthentication (new users will be denied access):

```
[edit protocols dot1x authenticator]
user@switch# set interface ge-0/0/1 server-fail use-cache
```

- Configure an interface that receives an EAPOL access-reject message from the authentication server to move end devices attempting LAN access on the interface to a specified VLAN already configured on the switch (in this case, the VLAN name is **vlan-sf**):

```
[edit protocols dot1x authenticator]
user@switch# set interface ge-0/0/1 server-reject-vlan vlan-sf
```

Related Documentation

- Example: Configuring 802.1X Authentication Options When the RADIUS Server is Unavailable to a J-EX Series Switch on page 2271
- Configuring 802.1X Authentication (J-Web Procedure) on page 2332
- Configuring 802.1X Interface Settings (CLI Procedure) on page 2331
- Monitoring 802.1X Authentication on page 2355
- Understanding Server Fail Fallback and Authentication on J-EX Series Switches on page 2258

Configuring 802.1X RADIUS Accounting (CLI Procedure)

RADIUS accounting permits statistical data about users logging onto or off a LAN to be collected and sent to a RADIUS accounting server. The statistical data gathered can be used for general network monitoring, to analyze and track usage patterns, or to bill a user based upon the amount of time or type of services accessed.

To configure basic RADIUS accounting using the CLI:

1. Specify the accounting servers to which the switch will forward accounting statistics:

```
[edit access]
user@switch# set profile profile1 radius accounting-server [122.69.1.250 122.69.1.252]
```

2. Define the RADIUS accounting servers:

```
[edit access]
user@switch# set radius-server 122.69.1.250 secret juniper

user@switch# set radius-server 122.69.1.252 secret juniper1
```

3. Enable accounting for an access profile:

```
[edit access]
user@switch# set profile profile1 accounting
```

4. Configure the RADIUS servers to use while sending accounting messages and updates:

```
[edit access]
user@switch# set profile profile1 accounting order radius none
```

5. Configure the statistics to be collected on the switch and forwarded to the accounting server:

```
[edit access]
user@switch# set profile profile1 accounting order accounting-stop-on-access-deny
user@switch# set profile profile1 accounting order accounting-stop-on-failure
```

6. Display accounting statistics collected on the switch:

```
user@switch> show network-access aaa statistics accounting

Accounting module statistics
  Requests received: 1
  Accounting Response failures: 0
  Accounting Response Success: 1
  Requests timedout: 0
```

7. Open an accounting log on the RADIUS accounting server using the server's address, and view accounting statistics:

```
[root@freeradius]# cd /usr/local/var/log/radius/radacct/122.69.1.250
[root@freeradius 122.69.1.250]# ls


detail-20071214


[root@freeradius 122.69.1.250]# vi details-20071214
```

```
User-Name = "000347e1bab9"
NAS-Port = 67
Acct-Status-Type = Stop
Acct-Session-Id = "802.1x811912"
Acct-Input-Octets = 17454
Acct-Output-Octets = 4245
Acct-Session-Time = 1221041249
Acct-Input-Packets = 72
Acct-Output-Packets = 53
Acct-Terminate-Cause = Lost-Carrier
Acct-Input-Gigawords = 0
Acct-Output-Gigawords = 0
Called-Station-Id = "00-19-e2-50-52-60"
Calling-Station-Id = "00-03-47-e1-ba-b9"
Event-Timestamp = "Sep 10 2008 16:52:39 PDT"
NAS-Identifier = "esp48t-1b-01"
NAS-Port-Type = Virtual

User-Name = "000347e1bab9"
NAS-Port = 67
Acct-Status-Type = Start
Acct-Session-Id = "802.1x811219"
Called-Station-Id = "00-19-e2-50-52-60"
Calling-Station-Id = "00-03-47-e1-ba-b9"
Event-Timestamp = "Sep 10 2008 18:58:52 PDT"
NAS-Identifier = "esp48t-1b-01"
NAS-Port-Type = Virtual
```

Related Documentation

- Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch on page 2267
- Understanding 802.1X and RADIUS Accounting on J-EX Series Switches on page 2260

Filtering 802.1X Supplicants Using RADIUS Server Attributes

There are two ways to configure the RADIUS server with port firewall filters:

- Include a match statement and corresponding action in the **Juniper-Firewall-Filter** attribute. The **Juniper-Firewall-Filter** attribute is a vendor-specific attribute (VSA) in the Juniper dictionary on the RADIUS server. Use this attribute to configure simple filter conditions for authenticated users. Nothing needs to be configured on the switch; all of the configuration is on the RADIUS server.
- Apply a local firewall filter to users authenticated through the RADIUS server. Use this method for more complex filters. The firewall filter must be configured on each switch.

This example describes using FreeRADIUS software to configure VSAs. For specifics on configuring your server, consult the AAA documentation that was included with your server.

This topic includes the following tasks:

1. Configuring Match Statements on the RADIUS Server on page 2341
2. Applying a Port Firewall Filter from the RADIUS Server on page 2343

Configuring Match Statements on the RADIUS Server

You can configure simple filter conditions using the **Juniper-Switching-Filter** attribute in the Juniper dictionary on the RADIUS server. These filters are then sent to a switch whenever a new user is authenticated successfully. The filters are created and applied on all J-EX Series switches that authenticate users through that RADIUS server without the need to configure anything on each individual switch.

To configure the **Juniper-Switching-Filter** attribute, enter one or more match conditions and a resulting action using the CLI for the RADIUS server. Enter the match statement plus an action statement enclosed within quotes (" ") using the following syntax:

```
match <destination-mac mac-address> <source-vlan vlan-name> <source-dot1q-tag
  tag> <destination-ip ip-address> <ip-protocol protocol-id> <source-port port>
  <destination-port port>
}
action [allow | deny] <forwarding-class class-of-service> <loss-priority (low | medium |
  high)>
}
```

See “VSA Match Conditions and Actions for J-EX Series Switches” on page 2348 for definitions of match statement options.

To configure match conditions on the RADIUS server:

1. Verify that the Juniper dictionary is loaded on your RADIUS server and includes the filtering attribute **Juniper-Switching-Filter**, attribute ID 48:

   ```
   [root@freeradius]# cat /usr/local/share/freeradius/dictionary.juniper


   #  dictionary.juniper
   #
   # Version:      $Id: dictionary.juniper,v 1.2.6.1 2005/11/30 22:17:25
   aland Exp
   $
   #  VENDOR          Juniper                          2636
   BEGIN-VENDOR    Juniper
   ATTRIBUTE       Juniper-Local-User-Name               1       string
   ATTRIBUTE       Juniper-Allow-Commands                2       string
   ATTRIBUTE       Juniper-Deny-Commands                 3       string
   ATTRIBUTE       Juniper-Allow-Configuration           4       string
   ATTRIBUTE       Juniper-Deny-Configuration            5       string
   ATTRIBUTE       Juniper-Switching-Filter              48      string
   <–
   ```

2. Enter the match conditions and actions. For example:

   - To deny authentication based on the 802.1Q tag (here, the 802.1Q tag is **10**):

     ```
     [root@freeradius]#
     cd /usr/local/etc/raddb
     vi users
     ```

     For each relevant user, add the **Juniper-Switching-Filter** attribute:

     ```
     Juniper-Switching-Filter = "match source-dot1q-tag 10 action deny"
     ```

- To deny access based on a destination IP address:

  ```
  [root@freeradius]# cd /usr/local/etc/raddb
  vi users
  ```

  For each relevant user, add the **Juniper-Switching-Filter** attribute:

  ```
  Juniper-Switching-Filter = “match destination-ip 192.168.1.0/31 action deny”
  ```

- To set the packet loss priority (PLP) to **high** based on a destination MAC address and the IP protocol:

  ```
  [root@freeradius]# cd /usr/local/etc/raddb
  vi users
  ```

  For each relevant user, add the **Juniper-Switching-Filter** attribute:

  ```
  Juniper-Switching-Filter = "match destination-mac 00:04:0f:fd:ac:fe, ip-protocol 2,
  forwarding-class high, action loss-priority high"
  ```

  NOTE: For the forwarding-class option to be applied, the forwarding class must be configured on the switch. If it is not configured on the switch, this option is ignored. You must specify both the forwarding class and the packet loss priority.

3. Stop and restart the RADIUS process to activate the configuration.

Applying a Port Firewall Filter from the RADIUS Server

You can apply a firewall filter to user policies on the RADIUS server. The RADIUS server can then specify the firewall filters that are to be applied to each user that requests to authenticate. Use this method when the firewall filter has more extensive conditions or you want to use different conditions for the same filter on different switches. The firewall filters must be configured on each switch.

For more information about firewall filters, see “Firewall Filters for J-EX Series Switches Overview” on page 2721.

To apply a port firewall filter centrally from the RADIUS server:

NOTE: If port firewall filters are also configured locally for the interface, then VSAs take precedence if they conflict with the filters. If the VSAs and the local port firewall filters do not conflict, they are merged.

1. Create the firewall filter on the local switch. In this example, the filter is called **filter1**.
2. Open the users file on the RADIUS server:

   ```
   [root@freeradius]#
   cd /usr/local/pool/raddb
   vi users
   ```

3. For each relevant user, add the filter (here, the filter ID is **filter1**):

   ```
   Filter-Id = "filter1"
   ```

   

   NOTE: Multiple filters are not supported on a single interface. However, you can support multiple filters for multiple users that are connected to the switch on the same interface by configuring a single filter with policies for each of those users.

4. Stop and restart the RADIUS process to activate the configuration.

Related Documentation

- Example: Applying a Firewall Filter to 802.1X-Authenticated Supplicants Using RADIUS Server Attributes on a J-EX Series Switch on page 2296
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Configuring 802.1X Interface Settings (CLI Procedure) on page 2331
- Understanding 802.1X and VSAs on J-EX Series Switches on page 2266

Configuring LLDP (CLI Procedure)

J-EX Series switches use Link Layer Discovery Protocol (LLDP) and Link Layer Discovery Protocol–Media Endpoint Discovery (LLDP-MED) to learn and distribute device information on network links. The information allows the switch to quickly identify a variety of devices, resulting in a LAN that interoperates smoothly and efficiently.

This topic describes:

Enabling LLDP on Interfaces

LLDP is enabled on all interfaces by default. If it is disabled, you can enable LLDP by configuring it on all interfaces or specific interfaces.

To configure LLDP on all interfaces or on a specific interface:

```
[edit protocols lldp]
user@switch# set interface all
```

Configuring for Fast Start

You can specify the number of LLDP-MED advertisements sent from the switch in the first second after it has detected an LLDP-capable device. The default is 3; to set it to another value:

```
[edit protocols lldp]
user@switch# set fast-start 8
```

Adjusting LLDP Advertisement Settings

You can adjust the following settings for LLDP advertisements for troubleshooting or verification purposes. For normal operations, we recommend that you do not adjust these settings from the default values.

- To specify the frequency at which LLDP advertisements are sent (in seconds):

  ```
  [edit protocols lldp]
  user@switch# set advertisement-interval 45
  ```

- To determine the length of time LLDP information is held before it is discarded (the multiplier value is used in combination with the **advertisement-interval** value):

  ```
  [edit protocols lldp]
  user@switch# set hold-multiplier 5
  ```

Adjusting SNMP Notification Settings of LLDP Changes

You can adjust the following settings for SNMP notifications of LLDP changes. If the values are not specified or the interval values are set to **0**, the notifications are disabled.

- To specify the frequency at which LLDP database changes are sent (in seconds):

```
[edit protocols lldp]
user@switch# set lldp-configuration-notification-interval 600
```

- To specify the frequency at which changes in topology global statistics are sent (in seconds):

```
[edit protocols lldp]
user@switch# set ptopo-configuration-trap-interval 600
```

- To specify the holding time (used in combination with the **ptopo-configuration-trap-interval** value) to determine the length of time that topology global statistics are held before they are is discarded (in seconds):

```
[edit protocols lldp]
user@switch# set ptopo-configuration-maximum-hold-time 2147483647
```

Specifying a Management Address for the LLDP Management TLV

You can configure an IP management address to be used in the LLDP Management type, length, and value (TLV).

To configure the management address:

```
[edit protocols lldp]
user@switch# set management-address 192.168.0.0
```

Related Documentation

- Configuring LLDP (J-Web Procedure) on page 2345
- Configuring LLDP-MED (CLI Procedure) on page 2346
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261

Configuring LLDP (J-Web Procedure)

Use the LLDP Configuration page to configure LLDP global and port settings for a J-EX Series switch on the J-Web interface.

To configure LLDP:

1. Select **Configure** > **Switching** > **LLDP**.

   The LLDP Configuration page displays LLDP Global Settings and Port Settings.

   The second half of the screen displays operational details for the selected port.

NOTE: After you make changes to the configuration in this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See "Using the Commit Options to Commit Configuration Changes (J-Web Procedure)" on page 334 for details about all commit options.

2. To modify LLDP Global Settings, click **Global Settings**.

   Enter information as described in Table 299 on page 2346.

3. To modify Port Settings, click **Edit** in the Port Settings section.

   Enter information as described in Table 300 on page 2346.

Table 299: Global Settings

Field	Function	Your Action
Advertising interval	Specifies the frequency of outbound LLDP advertisements. You can increase or decrease this interval.	Type the number of seconds.
Hold multiplier	Specifies the multiplier factor to be used by an LLDP-enabled switch to calculate the time-to-live (TTL) value for the LLDP advertisements it generates and transmits to LLDP neighbors.	Type the required number in the field.
Fast start count	Specifies the number of LLDP advertisements sent in the first second after the device connects. The default is 3. Increasing this number results in the port initially advertising LLDP–MED at a faster rate for a limited time.	Type the Fast start count.

Table 300: Edit Port Settings

Field	Function	Your Action
LLDP Status	Specifies whether LLDP has been enabled on the port.	Select one: **Enabled**, **Disabled**, or **None**.
LLDP-MED Status	Specifies whether LLDP–MED has been enabled on the port.	Select **Enable** from the list.

Related Documentation

- Configuring LLDP (CLI Procedure) on page 2344
- Configuring LLDP-MED (CLI Procedure) on page 2346
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261

Configuring LLDP-MED (CLI Procedure)

Link Layer Discovery Protocol–Media Endpoint Discovery (LLDP-MED) is an extension of LLDP. The J-EX Series switch uses LLDP-MED to support device discovery of VoIP telephones and to create location databases for these telephone locations.

LLDP-MED is turned on by default on J-EX Series switches.

This topic describes:

Enabling LLDP-MED on Interfaces

LLDP-MED is enabled on all interfaces by default. If it is disabled, you can enable LLDP-MED by configuring it on all interfaces or on specific interfaces.

To configure LLDP-MED on all interfaces or on a specific interface:

```
[edit protocols lldp-med]
user@switch# set interface ge-0/0/2.0
```

Configuring Location Information Advertised by the Switch

You can configure the location information that is advertised from the switch to the LLDP-MED device. You can specify a civic-based location (geographic location) or a location based on an elin (emergency location identification string):

- To specify a location by geography:

```
[edit protocols lldp-med]

user@switch# set interface ge-0/0/2.0 location civic-based country-code US

user@switch# set interface ge-0/0/2.0 location civic-based ca-type 1 ca-value “El Dorado County”

user@switch# set interface ge-0/0/2.0 location civic-based ca-type 2 ca-value CA

user@switch# set interface ge-0/0/2.0 location civic-based ca-type 3 ca-value Somerset

user@switch# set interface ge-0/0/2.0 location civic-based ca-type 6 ca-value “Mount Aukum Road”

user@switch# set interface ge-0/0/2.0 location civic-based ca-type 19 ca-value 6450

user@switch# set interface ge-0/0/2.0 location civic-based ca-type 21 ca-value “Holiday Market”
```

- To specify a location using an **elin** string:

```
[edit protocols lldp-med]

user@switch# set interface ge-0/0/2.0 location elin 4085551212
```

Configuring for Fast Start

You can specify the number of LLDP-MED advertisements sent from the switch in the first second after it has detected an LLDP-MED device. The default is 3; to set it to another value:

```
[edit protocols lldp-med]
user@switch# set fast-start 6
```

Related Documentation

- Configuring LLDP (J-Web Procedure) on page 2345
- Example: Setting Up VoIP with 802.1X and LLDP-MED on a J-EX Series Switch on page 2302
- Configuring LLDP (CLI Procedure) on page 2344
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261

VSA Match Conditions and Actions for J-EX Series Switches

J-EX Series switches support the configuration of RADIUS server attributes specific to Juniper Networks. These attributes are known as vendor-specific attributes (VSAs). They are configured on RADIUS servers and work in combination with 802.1X authentication. Using VSAs, you can apply port firewall filter attributes as a subset of match conditions and actions sent from the RADIUS server to the switch as a result of 802.1X authentication success.

Each term in a VSA configured through the RADIUS server consists of *match conditions* and an *action*. Match conditions are the values or fields that the packet must contain. You can define single, multiple, or no match conditions. If no match conditions are specified for the term, the packet is accepted by default. The action is the action that the switch takes if a packet matches the match conditions for the specific term. Allowed actions are accept a packet or discard a packet.

The following guidelines apply when you specify match conditions and actions for VSAs:

- Both **match** and **action** statements are mandatory.
- Any or all options (separated by commas) may be included in each **match** and **action** statement.
- Fields separated by commas will be ANDed if they are of a different type. The same types cannot be repeated.
- For OR cases (for example, match **10.1.1.0/24** OR **11.1.1.0/24**), apply multiple VSAs to the 802.1X supplicant.
- In order for the **forwarding-class** option to be applied, the forwarding class must be configured on the switch. If it is not configured on the switch, this option is ignored.

Table 301 on page 2348 describes the match conditions you can specify when configuring a VSA using the **match** command on the RADIUS server. The string that defines a match condition is called a *match statement*.

Table 301: Match Conditions

Option	Description
destination-mac *mac-address*	Destination media access control (MAC) address of the packet.
source-vlan *source-vlan*	Name of the source VLAN.

Table 301: Match Conditions *(continued)*

Option	Description
source-dot1q-tag *tag*	Tag value in the dot1q header, in the range **0** through **4095**.
destination-ip *ip-address*	Address of the final destination node.
ip-protocol *protocol-id*	IPv4 protocol value. In place of the numeric value, you can specify one of the following text synonyms: **ah**, **egp (8)**, **esp (50**, **gre (47)**, **icmp (1)**, **igmp (2)**, **ipip (4)**, **ipv6 (41)**, **ospf (89)**, **pim (103)**, **rsvp (46)**, **tcp (6)**, or **udp (17)**
source-port *port*	TCP or User Datagram Protocol (UDP) source port field. Normally, you specify this match statement in conjunction with the **ip-protocol** match statement to determine which protocol is being used on the port. In place of the numeric field, you can specify one of the text options listed under **destination-port**.
destination-port *port*	TCP or UDP destination port field. Normally, you specify this match in conjunction with the **ip-protocol** match statement to determine which protocol is being used on the port. In place of the numeric value, you can specify one of the following text synonyms (the port numbers are also listed): **afs (1483)**, **bgp (179)**, **biff (512)**, **bootpc (68)**, **bootps (67)**, **cvspserver (2401)**, **cmd (514)**, **dhcp (67)**, **domain (53)**, **eklogin (2105)**, **ekshell (2106)**, **exec (512)**, **finger (79)**, **ftp (21)**, **ftp-data (20)**, **http (80)**, **https (443)**, **ident (113)**, **imap (143)**, **kerberos-sec (88)**, **klogin (543)**, **kpasswd (761)**, **krb-prop (754)**, **krbupdate (760)**, **kshell (544)**, **ldap (389)**, **login (513)**, **mobileip-agent (434)**, **mobilip-mn (435)**, **msdp (639)**, **netbios-dgm (138)**, **netbios-ns (137)**, **netbios-ssn (139)**, **nfsd (2049)**, **nntp (119)**, **ntalk (518)**, **ntp (123)**, **pop3 (110)**, **pptp (1723)**, **printer (515)**, **radacct (1813)**, **radius (1812)**, **rip (520)**, **rkinit (2108)**, **smtp (25)**, **snmp (161)**, **snmptrap (162)**, **snpp (444)**, **socks (1080)**, **ssh (22)**, **sunrpc (111)**, **syslog (514)**, **telnet (23)**, **tacacs-ds (65)**, **talk (517)**, **tftp (69)**, **timed (525)**, **who (513)**, **xdmcp (177)**, **zephyr-clt (2103)**, **zephyr-hm (2104)**

When you define one or more terms that specify the filtering criteria, you also define the action to take if the packet matches all criteria. Table 302 on page 2349 shows the actions that you can specify in a term.

Table 302: Actions for VSAs

Option	Description
(**allow** | **deny**)	Accept a packet or discard a packet silently without sending an Internet Control Message Protocol (ICMP) message.
forwarding-class *class-of-service*	(Optional) Classify the packet in one of the following forwarding classes: • **assured-forwarding** • **best-effort** • **expedited-forwarding** • **network-control**

Table 302: Actions for VSAs *(continued)*

Option	Description
loss-priority (**low** | **medium** | **high**)	(Optional) Set the packet loss priority (PLP) to **low**, **medium**, or **high**. Specify both the forwarding class and loss priority.

Related Documentation

- Filtering 802.1X Supplicants Using RADIUS Server Attributes on page 2340
- Understanding 802.1X and VSAs on J-EX Series Switches on page 2266

Configuring Captive Portal Authentication (CLI Procedure)

Configure captive portal authentication (hereafter referred to as captive portal) on a J-EX Series switch so that users connected to the switch are authenticated before being allowed to access the network. When the user requests a webpage, a login page is displayed that requires the user to input a username and password. Upon successful authentication, the user is allowed to continue with the original page request and subsequent access to the network.

Before you begin, be sure you have:

- Performed basic bridging and VLAN configuration on the switch. See "Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch" on page 1063.
- Generated an SSL certificate and installed it on the switch. See "Generating SSL Certificates to Be Used for Secure Web Access" on page 398.
- Configured basic access between the J-EX Series switch and the RADIUS server. See "Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch" on page 2267.
- Designed your captive portal login page. See "Designing a Captive Portal Authentication Login Page on a J-EX Series Switch" on page 2351.

This topic includes the following tasks:

Configuring Secure Access for Captive Portal

To configure secure access for captive portal:

1. Associate the security certificate with the Web server and enable HTTPS on the switch:

```
[edit]
user@switch# set system services web-management https local-certificate
my-signed-cert
```

NOTE: You can enable HTTP instead of HTTPS, but we recommend HTTPS for security purposes.

2. Configure captive portal to use HTTPS:

```
[edit]
user@switch# set services captive-portal secure-authentication https
```

Enabling an Interface for Captive Portal

To enable an interface for use with captive portal authentication:

```
[edit]
user@switch# set services captive-portal interface ge-0/0/10
```

Configuring Bypass of Captive Portal Authentication

You can allow specific clients to bypass captive portal authentication:

```
[edit]
user@switch# set ethernet-switching-options authentication-whitelist 00:10:12:e0:28:22
```

NOTE: Optionally, you can use set ethernet-switching-options authentication-whitelist 00:10:12:e0:28:22 interface ge-0/0/10.0 to limit the scope to the interface.

If the MAC address of the client that you want to configure for authentication bypass has already been learned on the interface, you must clear it using the **clear captive-portal interface *interface-name*** before adding it to the whitelist. Otherwise the new entry for the MAC address will not be added to the ethernet switching table and the authentication bypass will not be allowed.

Related Documentation

- Example: Setting Up Captive Portal Authentication on a J-EX Series Switch on page 2323
- Understanding Captive Portal Authentication

Designing a Captive Portal Authentication Login Page on a J-EX Series Switch

You can set up captive portal authentication on your switch to redirect all Web browser requests to a login page that requires the user to input a username and password before they are allowed access. Upon successful authentication, the user is allowed access to the network and to continue to the original page requested.

The Junos OS provides a customizable template for the captive portal window that allows you to easily design and modify the look of the captive portal login page. You can modify the design elements in the template to change the look of your captive portal login page and to add instructions or information to the page. You can also modify any of the design elements of an existing captive portal login page.

Figure 59 on page 2352 shows an example of a captive portal login page:

Figure 59: Example of a Captive Portal Login Page

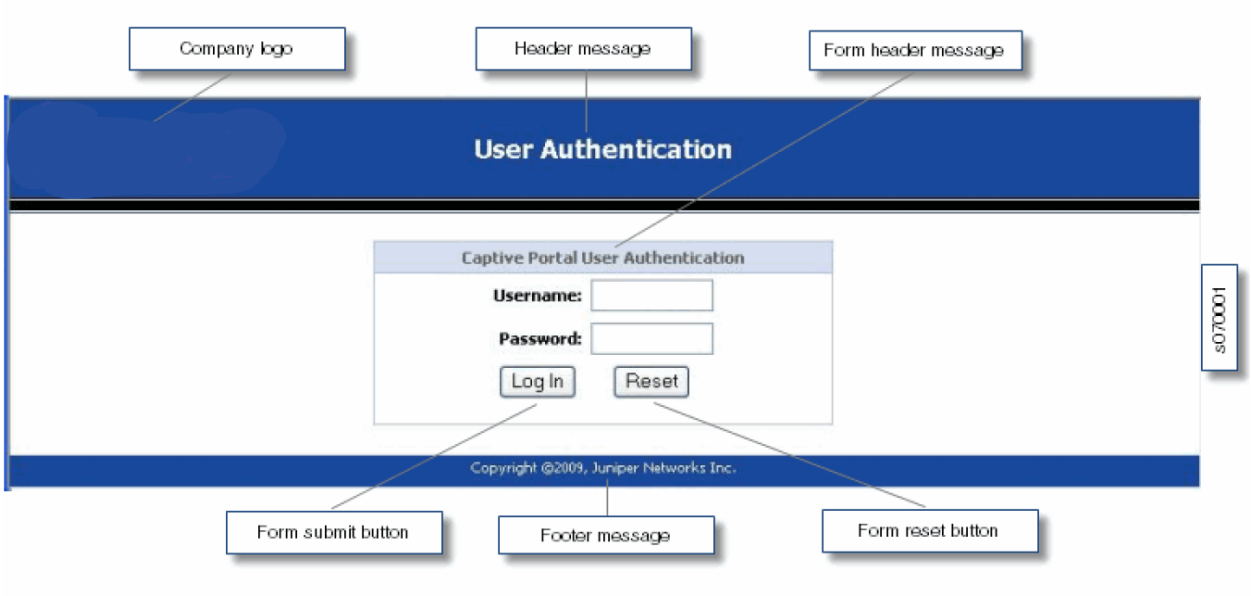

Table 303 on page 2352 summarizes the configurable elements of a captive portal login page.

Table 303: Configurable Elements of a Captive Portal Login Page

Element	CLI Statement	Description
Banner message	**banner-message** *text-string*	The first screen displayed before the captive portal login page is displayed (not shown). The page header says "**Terms and Conditions of Use: Please read the following terms of use and disclaimers carefully before using this network.**" The configurable banner message appears in the body of the page. The default text is "**Terms and Conditions**." A button labeled **Agree** gives the user access to the captive portal login page.
Footer background color	**footer-bgcolor** *hex-color*	The HTML hexadecimal code for the background color of the captive portal login page footer.
Footer message	**footer-message** *text-string*	For example, you can include copyright information and links to additional information such as help instructions, legal notices, or a privacy policy.
Form header background color	**form-header-bgcolor** *hex-color*	The HTML hexadecimal code for the background color of the header bar across the top of the form area of the captive portal login page.
Form header message	**form-header-message** *text-string*	Text displayed in the header bar across the top of the form area of the captive portal login page. For example, **Welcome to My Cafe**. The default text is **Captive Portal User Authentication**.
Form reset button label	**form-reset-label** *label-name*	Label appearing in the button that the user can select to clear the username and password fields on the form, for example, **Reset** or **Clear**.

Table 303: Configurable Elements of a Captive Portal Login Page *(continued)*

Element	CLI Statement	Description
Form submit button label	**form-submit-label** ***label-name***	Label appearing in the button that user selects to submit their login information—for example, **Log In** or **OK**.
Header background color	**header-bgcolor** ***hex-color***	The HTML hexadecimal code for the background color of the captive portal login page header.
Header logo	**header-logo** ***filename***	Filename of the file containing the image of the logo that you want to appear at the top of the captive portal login page. The image file can be in GIF, JPEG, or PNG format. You can upload a logo image file to the switch. Copy the logo to the /var/tmp directory on the switch (during the commit the files are saved to persistent locations).
Header message	**header-message** ***text-string***	Text displayed in the page header. The default text is **User Authentication**.
Post-authentication URL	**post-authentication-url** ***url***	URL to which the users are directed upon successful authentication. The default is to redirect users to the page they had originally requested.

To design the captive portal login page:

1. (Optional) Upload your logo image file to the switch:

```
user@switch> file copy
ftp://username:prompt@ftp.hostname.net/var/tmp/my-logo.jpeg
```

2. Configure the custom options to specify the background colors and text displayed in the captive portal page:

```
[edit system services captive-portal]
user@switch# set custom-options header-bgcolor #006600
set custom-options header-message “Welcome to Our Network”
set custom-options banner-message “Please enter your username and password:”
set custom-options footer-message “Copyright ©2009, Our Network”
```

NOTE: For the custom options that you do not specify, the value is taken from the standard template.

Related Documentation

- Example: Setting Up Captive Portal Authentication on a J-EX Series Switch on page 2323

CHAPTER 84

Verifying 802.1X and MAC RADIUS Authentication

Monitoring 802.1X Authentication

Purpose Use the monitoring feature to display details of authenticated users and users who have failed authentication.

Action To display authentication details in the J-Web interface, select **Monitoring > Security > 802.1X**.

To display authentication details in the CLI, enter the following commands:

- **show dot1x interface detail | display xml**
- **show dot1x interface detail <*interface*> | display xml**
- **show dot1x auth-failed-users**

Meaning The details displayed include:

- A list of authenticated users.
- The total number of users connected.
- A list of users who have failed authentication

You can also specify an interface for which the details must be displayed.

Related Documentation

- Configuring 802.1X Authentication (J-Web Procedure) on page 2332
- Configuring 802.1X Interface Settings (CLI Procedure) on page 2331
- Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on a J-EX Series Switch on page 2290

Verifying 802.1X Authentication

Purpose Verify that supplicants are being authenticated on an interface on a J-EX Series switch with the interface configured for 802.1X authentication, and display the method of authentication being used.

Action Display detailed information about an interface configured for 802.1X (here, the interface is **ge-0/0/16**):

```
user@switch> show dot1x interface ge-0/0/16.0 detail
ge-0/0/16.0
  Role: Authenticator
  Administrative state: Auto
  Supplicant mode: Single
  Number of retries: 3
  Quiet period: 60 seconds
  Transmit period: 30 seconds
  Mac Radius: Enabled
  Mac Radius Strict: Disabled
  Reauthentication: Enabled Reauthentication interval: 40 seconds
  Supplicant timeout: 30 seconds
  Server timeout: 30 seconds
  Maximum EAPOL requests: 1
  Guest VLAN member: <not configured>
  Number of connected supplicants: 1
    Supplicant: user5, 00:30:48:8C:66:BD
      Operational state: Authenticated
      Authentication method: Radius
      Authenticated VLAN: v200
      Reauthentication due in 17 seconds
```

Meaning The sample output from the **show dot1x interface detail** command shows that the **Number of connected supplicants** is 1. The supplicant that was authenticated and is now connected to the LAN is known as **user5** on the RADIUS server and has the MAC address **00:30:48:8C:66:BD**. The supplicant was authenticated by means of the 802.1X authentication method called **Radius** authentication. When the **Radius** authentication method is used, the supplicant is configured on the RADIUS server, the RADIUS server communicates this to the switch, and the switch opens LAN access on the interface to which the supplicant is connected. The sample output also shows that the supplicant is connected to VLAN **v200**.

Other 802.1X authentication methods supported on J-EX Series switches in addition to the **RADIUS** method are:

- **Guest VLAN**—A nonresponsive host is granted Guest-VLAN access.
- **MAC Radius**—A nonresponsive host is authenticated based on its MAC address. The MAC address is configured as permitted on the RADIUS server, the RADIUS server lets the switch know that the MAC address is a permitted address, and the switch opens LAN access to the nonresponsive host on the interface to which it is connected.
- **Server-fail deny**—If the RADIUS servers time out, all supplicants are denied access to the LAN, preventing traffic from flowing from the supplicant through the interface. This is the default.

- **Server-fail permit**—When the RADIUS server is unavailable, a supplicant is still permitted access to the LAN as if the supplicant had been successfully authenticated by the RADIUS server.
- **Server-fail use-cache**—If the RADIUS servers time out during reauthentication, previously authenticated supplicants are granted access, but new supplicants are denied LAN access.
- **Server-fail VLAN**—A supplicant is configured to be moved to a specified VLAN if the RADIUS server is unavailable to reauthenticate the supplicant. (The VLAN must already exist on the switch.)

Related Documentation

- Configuring 802.1X Interface Settings (CLI Procedure) on page 2331
- Configuring 802.1X Authentication (J-Web Procedure) on page 2332
- Configuring MAC RADIUS Authentication (CLI Procedure) on page 2335
- Configuring Server Fail Fallback (CLI Procedure) on page 2337

CHAPTER 85

Configuration Statements for Access Control

[edit access] Configuration Statement Hierarchy

```
access {
  profile profile-name {
    accounting {
      order [ radius | none ];
      accounting-stop-on-access-deny;
      accounting-stop-on-failure;
    }
    authentication-order [ authentication-method ];
    radius {
      accounting-server [ server-address ];
      authentication-server [ server-address ];
    }
  }
}
```

Related Documentation

- Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch on page 2267
- Configuring 802.1X RADIUS Accounting (CLI Procedure) on page 2339

[edit ethernet-switching-options] Configuration Statement Hierarchy

```
ethernet-switching-options {
  analyzer {
    name {
      loss-priority priority;
      ratio number;
      input {
        ingress {
          interface (all | interface-name);
          vlan (vlan-id | vlan-name);
        }
```

```
        egress {
          interface (all | interface-name);
        }
      output {
        interface interface-name;
        vlan (vlan-id | vlan-name);
      }
    }
  }
  bpdu-block {
    disable-timeout timeout;
    interface (all | [interface-name]);
  }
  dot1q-tunneling {
    ether-type (0x8100 | 0x88a8 | 0x9100);
  }
  interfaces interface-name {
    no-mac-learning;
  }
  mac-notification {
    notification-interval seconds;
  }
  mac-table-aging-time seconds;
  port-error-disable {
    disable-timeout timeout;
  }
  redundant-trunk-group {
    group-name name {
      interface interface-name <primary>;
    }
  }
  secure-access-port {
    dhcp-snooping-file {
      location local_pathname | remote_URL;
      timeout seconds;
      write-interval seconds;
    }
    interface (all | interface-name) {
      allowed-mac {
        mac-address-list;
      }
      (dhcp-trusted | no-dhcp-trusted );
      mac-limit limit action action;
      no-allowed-mac-log;
      static-ip  ip-address {
        vlan  vlan-name;
        mac mac-address;
      }
    }
    vlan (all | vlan-name) {
      (arp-inspection | no-arp-inspection );
      dhcp-option82 {
        circuit-id {
          prefix hostname;
          use-interface-description;
          use-vlan-id;
```

```
            }
            remote-id {
              prefix hostname | mac | none;
              use-interface-description;
              use-string string;
            }
            vendor-id [string];
          }
          (examine-dhcp | no-examine-dhcp );
          (ip-source-guard | no-ip-source-guard);
          mac-move-limit limit action action;
        }
      }
      storm-control {
        action-shutdown;
        interface (all | interface-name) {
          bandwidth bandwidth;
          no-broadcast;
          no-unknown-unicast;
        }
      }
      traceoptions {
        file filename <files number> <no-stamp> <replace> <size size> <world-readable |
          no-world-readable>;
        flag flag <disable>;
      }
      unknown-unicast-forwarding {
        vlan (all | vlan-name) {
          interface interface-name;
        }
      }
      voip {
        interface (all | [interface-name | access-ports]) {
          vlan vlan-name ;
          forwarding-class <assured-forwarding | best-effort | expedited-forwarding |
            network-control>;
        }
      }
    }
```

Related Documentation

- Understanding Port Mirroring on J-EX Series Switches on page 3245
- Port Security for J-EX Series Switches Overview on page 2545
- Understanding BPDU Protection for STP, RSTP, and MSTP on J-EX Series Switches on page 1278
- Understanding Redundant Trunk Links on J-EX Series Switches on page 1049
- Understanding Storm Control on J-EX Series Switches on page 2511
- Understanding 802.1X and VoIP on J-EX Series Switches on page 2263
- Understanding Q-in-Q Tunneling on J-EX Series Switches on page 1051
- Understanding Unknown Unicast Forwarding on J-EX Series Switches on page 2512
- Understanding MAC Notification on J-EX Series Switches on page 1060

[edit protocols] Configuration Statement Hierarchy

```
protocols {
  connections {
    remote-interface-switch connection-name {
      interface interface-name.unit-number;
      transmit-lsp label-switched-path;
      receive-lsp label-switched-path;
    }
  }
  dot1x {
    authenticator {
      authentication-profile-name profile-name;
      interface (all | [ interface-names ]) {
        disable;
        guest-vlan ( vlan-id | vlan-name);
        mac-radius <restrict>;
        maximum-requests number;
        no-reauthentication;
        quiet-period seconds;
        reauthentication {
          interval seconds;
        }
        retries number;
        server-fail (deny | permit | use-cache | vlan-id | vlan-name);
        server-reject-vlan (vlan-id | vlan-name);
        server-timeout seconds;
        supplicant (multiple | single | single-secure);
        supplicant-timeout seconds;
        transmit-period seconds;
      }
      static mac-address {
        interface interface-name;
        vlan-assignment (vlan-id |vlan-name);
      }
  }
  gvrp {
    <enable | disable>;
    interface (all | [interface-name]) {
      disable;
    }
    join-timer millseconds;
    leave-timer milliseconds;
    leaveall-timer milliseconds;
  }
  igmp-snooping {
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>
        <match regex>;
      flag flag (detail | disable | receive | send);
    }
    vlan (vlan-id | vlan-number) {
      data-forwarding {
        source {
```

```
            groups group-prefix;
          }
          receiver {
            source-vlans vlan-list;
            install ;
          }
        }
        disable {
          interface interface-name
        }
        immediate-leave;
        interface interface-name {
          group-limit limit;
          multicast-router-interface;
          static {
            group ip-address;
          }
        }
        proxy ;
        query-interval seconds;
        query-last-member-interval seconds;
        query-response-interval seconds;
        robust-count number;
      }
    }
    lldp {
      disable;
      advertisement-interval seconds;
      hold-multiplier number;
      interface (all | interface-name) {
        disable;
      }
      lldp-configuration-notification-interval seconds;
      management-address ip-management-address;
      ptopo-configuration-maximum-hold-time seconds;
      ptopo-configuration-trap-interval seconds;
      traceoptions {
        file filename <files number> <size size> <world-readable | no-world-readable>
          <match regex>;
        flag flag (detail | disable | receive | send);
      }
    }
    lldp-med {
      disable;
      fast-start number;
      interface (all | interface-name) {
        disable;
        location {
          elin number;
          civic-based {
            what number;
            country-code code;
            ca-type {
              number {
                ca-value value;
              }
```

```
        }
      }
    }
  }
}
mpls {
  interface ( all | interface-name );
  label-switched-path lsp-name to remote-provider-edge-switch;
  path destination {
    <address | hostname> <strict | loose>
  }
mstp {
  disable;
  bpdu-block-on-edge;
  bridge-priority priority;
  configuration-name name;
  forward-delay seconds;
  hello-time seconds;
  interface (all | interface-name) {
    disable;
    bpdu-timeout-action {
      block;
      alarm;
    }
    cost  cost;
    edge;
    mode mode;
    no-root-port;
    priority priority;
  }
  max-age seconds;
  max-hops hops;
  msti msti-id {
    vlan (vlan-id | vlan-name);
    interface interface-name {
      disable;
      cost  cost;
      edge;
      mode mode;
      priority priority;
    }
  }
  revision-level revision-level;
  traceoptions {
    file filename <files number > <size size> <no-stamp | world-readable |
      no-world-readable>;
    flag flag;
  }
}
mvrp {
  disable
  interface (all | interface-name) {
    disable;
    join-timer milliseconds;
    leave-timer milliseconds;
    leaveall-timer milliseconds;
```

```
    registration (forbidden | normal);
  }
  no-dynamic-vlan;
  traceoptions {
    file filename <files number > <size size> <no-stamp | world-readable |
      no-world-readable>;
    flag flag;
  }
}
oam {
  ethernet{
    connectivity-fault-management {
      action-profile profile-name {
        default-actions {
          interface-down;
        }
      }
      linktrace {
        age (30m | 10m | 1m | 30s | 10s);
        path-database-size path-database-size;
      }
      maintenance-domain domain-name {
        level number;
        mip-half-function (none | default |explicit);
        name-format (character-string | none | dns | mac+2oct);
        maintenance-association ma-name {
          continuity-check {
            hold-interval minutes;
            interval (10m | 10s | 1m | 1s| 100ms);
            loss-threshold number;
          }
          mep mep-id {
            auto-discovery;
            direction down;
            interface interface-name;
            remote-mep mep-id {
              action-profile profile-name;
            }
          }
        }
      }
    }
    link-fault-management {
      action-profile profile-name;
        action {
          syslog;
          link-down;
        }
        event {
          link-adjacency-loss;
          link-event-rate;
          frame-error count;
          frame-period count;
          frame-period-summary count;
          symbol-period count;
        }
```

```
            interface interface-name {
              link-discovery (active | passive);
              pdu-interval interval;
              event-thresholds threshold-value;
              remote-loopback;
              event-thresholds {
                frame-errorcount;
                frame-period count;
                frame-period-summary count;
                symbol-period count;
              }
            }
            negotiation-options {
              allow-remote-loopback;
              no-allow-link-events;
            }
        }
    }
}
rstp {
    disable;
    bpdu-block-on-edge;
    bridge-priority priority;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
        disable;
        bpdu-timeout-action {
            block;
            alarm;
        }
        cost  cost;
        edge;
        mode mode;
        no-root-port;
        priority priority;
    }
    max-age seconds;
    }
    traceoptions {
        file filename <files number > <size size> <no-stamp | world-readable |
            no-world-readable>;
        flag flag;
    }
}
sflow {
    agent-id;
    collector {
        ip-address;
        udp-port port-number;
    }
    disable;
    interfaces interface-name {
        disable;
        polling-interval seconds;
        sample-rate {
```

```
        egress number;
        ingress number;
      }
    }
    polling-interval seconds;
    sample-rate {
      egress number;
      ingress number;
    }
    source-ip;
  }
  stp {
    disable;
    bridge-priority priority;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
      disable;
      bpdu-timeout-action {
        block;
        log;
      }
      cost  cost;
      edge;
      mode mode;
      no-root-port;
      priority priority;
    }
    max-age seconds;
  }
    traceoptions {
      file filename <files number > <size size> <no-stamp | world-readable |
        no-world-readable>;
      flag flag;
    }
  vstp {
    bpdu-block-on-edge;
    disable;
    force-version stp;
    vlan (all | vlan-id | vlan-name) {
      bridge-priority priority;
      forward-delay seconds;
      hello-time seconds;
      interface (all | interface-name) {
        bpdu-timeout-action {
          alarm;
          block;
        }
        cost  cost;
        disable;
        edge;
        mode mode;
        no-root-port;
        priority priority;
      }
      max-age seconds;
```

```
        traceoptions {
          file filename <files number > <size size> <no-stamp | world-readable |
            no-world-readable>;
          flag flag;
        }
      }
    }
  }
```

Related Documentation

- 802.1X for J-EX Series Switches Overview on page 2253
- Example: Configure Automatic VLAN Administration Using GVRP on page 1087
- Understanding MAC RADIUS Authentication on J-EX Series Switches
- Understanding Server Fail Fallback and 802.1X Authentication on J-EX Series Switches on page 2258
- IGMP Snooping on J-EX Series Switches Overview on page 2047
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261
- Understanding MSTP for J-EX Series Switches on page 1277
- Understanding Multiple VLAN Registration Protocol (MVRP) on J-EX Series Switches on page 1054
- Understanding Ethernet OAM Connectivity Fault Management for a J-EX Series Switch on page 3463
- Understanding Ethernet OAM Link Fault Management for a J-EX Series Switch on page 3427
- Understanding RSTP for J-EX Series Switches on page 1276
- Understanding STP for J-EX Series Switches on page 1275
- Understanding How to Use sFlow Technology for Network Monitoring on a J-EX Series Switch on page 3283
- Understanding VSTP for J-EX Series Switches on page 1281

access

Syntax

```
access {
  profile profile-name {
    authentication-order [ldap radius | none];
    accounting {
      order [radius | none];
      accounting-stop-on-access-deny;
      accounting-stop-on-failure;
    }
    radius {
      accounting-server [server-addresses];
      authentication-server [server-addresses];
    }
  }
}
```

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure authentication, authorization, and accounting (AAA) services.

The statements are explained separately.

Default Not enabled

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch on page 2267
- Configuring 802.1X RADIUS Accounting (CLI Procedure) on page 2339

accounting

Syntax

```
accounting {
  order radius | none;
    accounting-stop-on-access-deny;
    accounting-stop-on-failure;
  }
}
```

Hierarchy Level [edit **access profile** *profile-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the authentication order for authentication, authorization, and accounting (AAA) services.

Default Not enabled

Options **none**—Use no authentication for specified subscribers.

radius—Use RADIUS authentication for specified subscribers.

The remaining statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch on page 2267
- Configuring 802.1X RADIUS Accounting (CLI Procedure) on page 2339
- Understanding 802.1X and RADIUS Accounting on J-EX Series Switches on page 2260

accounting (Access Profile)

Syntax

```
accounting {
  accounting-stop-on-access-deny;
  accounting-stop-on-failure;
  coa-immediate-update;
  immediate-update;
  order [ accounting-method ];
  statistics (time | volume-time);
  update-interval minutes;
}
```

Hierarchy Level [edit access profile *profile-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure RADIUS accounting parameters and enable RADIUS accounting for an access profile.

The remaining statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Authentication and Accounting Parameters for Subscriber Access
- Configuring How Accounting Statistics Are Collected for Subscriber Access

accounting

Syntax

```
accounting {
  events [ login change-log interactive-commands ];
  destination {
    radius {
      server {
        server-address {
          accounting-port port-number;
          secret password;
          source-address address;
          retry number;
          timeout seconds;
        }
      }
    }
    tacplus {
      server {
        server-address {
          port port-number;
          secret password;
          single-connection;
          timeout seconds;
        }
      }
    }
  }
}
```

Hierarchy Level [edit system]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure audit of TACACS+ or RADIUS authentication events, configuration changes, and interactive commands.

Options The remaining statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring RADIUS System Accounting
- Configuring TACACS+ System Accounting

accounting-port

Syntax accounting-port *port-number*;

Hierarchy Level [edit system accounting destination radius server *server-address*],
[edit system radius-server *server-address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the accounting port number on which to contact the RADIUS server.

Options *number*—Port number on which to contact the RADIUS server.
Default: 1813

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring RADIUS Authentication
- Configuring RADIUS System Accounting

accounting-server

Syntax accounting-server [*server-addresses*];

Hierarchy Level [edit access profile *profile-name* radius]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the Remote Authentication Dial-In User Service (RADIUS) server for authentication. To configure multiple RADIUS servers, include multiple server addresses. The servers are tried in order and in a round-robin fashion until a valid response is received from one of the servers or until all the configured retry limits are reached.

Default Not enabled

Options *server-addresses*—One or more addresses of RADIUS authentication servers.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- **show network-access aaa statistics authentication on page 2506**
- Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch on page 2267
- Understanding 802.1X and RADIUS Accounting on J-EX Series Switches on page 2260

accounting-session-id-format

Syntax accounting-session-id-format (decimal | description);

Hierarchy Level [edit access profile *profile-name* radius **options**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the format the router or switch uses to identify the accounting session.

Default **decimal**

Options **decimal**—Use the decimal format.

description—Use the generic format, in the form: **jnpr** *interface-specifier*:*subscriber-session-id*.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring RADIUS Server Options for Subscriber Access
- Configuring Authentication and Accounting Parameters for Subscriber Access

accounting-stop-on-access-deny

Syntax accounting-stop-on-access-deny;

Hierarchy Level [edit **access profile** *profile-name* **accounting**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configures the authentication order for authentication, authorization, and accounting (AAA) services to send an Acct-Stop message if the AAA server denies access to a supplicant.

Default Not enabled

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch on page 2267
- Configuring 802.1X RADIUS Accounting (CLI Procedure) on page 2339
- **show network-access aaa statistics authentication on page 2506**

accounting-stop-on-access-deny

Syntax accounting-stop-on-access-deny;

Hierarchy Level [edit access profile *profile-name* accounting]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure RADIUS accounting to send an Acct-Stop message when the AAA server refuses a client request for access.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring Authentication and Accounting Parameters for Subscriber Access

accounting-stop-on-failure

Syntax accounting-stop-on-failure;

Hierarchy Level [edit access profile *profile-name* accounting]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure authentication order for authentication, authorization, and accounting (AAA) services to send an Acct-Stop message if a supplicant fails AAA authorization, but the RADIUS server grants access. For example, a supplicant might fail AAA authentication due to an internal error such as a timeout.

Default Not enabled

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch on page 2267
- Configuring 802.1X RADIUS Accounting (CLI Procedure) on page 2339
- Understanding 802.1X and RADIUS Accounting on J-EX Series Switches on page 2260

accounting-stop-on-failure

Syntax accounting-stop-on-failure;

Hierarchy Level [edit access profile *profile-name* **accounting**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure RADIUS accounting to send an Acct-Stop message when client access fails AAA but the AAA server grants access.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring Authentication and Accounting Parameters for Subscriber Access

address

Syntax address *address-or-prefix*;

Hierarchy Level [edit access **address-pool** *pool-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the IP address or prefix value for clients.

Options ***address-or-prefix***—An address or prefix value.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring the Address Pool for L2TP Network Server IP Address Allocation

address-pool

Syntax

```
address-pool pool-name {
    address address-or-prefix;
    address-range <low lower-limit> <high upper-limit>;
}
```

Hierarchy Level [edit access]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Allocate IP addresses for clients.

Options ***pool-name***—Name assigned to an address pool.

The remaining statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring the Address Pool for L2TP Network Server IP Address Allocation

address-range

Syntax address-range <low *lower-limit*> <high *upper-limit*>;

Hierarchy Level [edit access **address-pool** *pool-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the address range.

Options

- **high** ***upper-limit***—Upper limit of an address range.
- **low** ***lower-limit***—Lower limit of an address range.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring the Address Pool for L2TP Network Server IP Address Allocation

advertisement-interval

Syntax advertisement-interval *seconds*;

Hierarchy Level [edit protocols lldp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For switches configured for Link Layer Discovery Protocol, configure the frequency at which LLDP advertisements are sent.

Default Disabled.

Options *seconds*—(Optional) The number of seconds.

Range: 5 through 32,768 seconds

Default: 30 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show lldp on page 2489**
- Configuring LLDP (CLI Procedure) on page 2344
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261

attributes

Syntax

```
attributes {
  exclude {
    accounting-authentic [ accounting-on | accounting-off ];
    accounting-delay-time [ accounting-on | accounting-off ];
    accounting-session-id [ access-request | accounting-on | accounting-off | accounting-stop
      ];
    accounting-terminate-cause [ accounting-off ];
    called-station-id [ access-request | accounting-start | accounting-stop ];
    calling-station-id [ access-request | accounting-start | accounting-stop ];
    class [ accounting-start | accounting-stop ];
    dhcp-gi-address [ access-request | accounting-start | accounting-stop ];
    dhcp-mac-address [ access-request | accounting-start | accounting-stop ];
    output-filter [ accounting-start | accounting-stop ];
    event-timestamp [ accounting-on | accounting-off | accounting-start | accounting-stop
      ];
    framed-ip-address [ accounting-start | accounting-stop ];
    framed-ip-netmask [ accounting-start | accounting-stop ];
    input-filter [ accounting-start | accounting-stop ];
    input-gigapackets [ accounting-stop ];
    input-gigawords [ accounting-stop ];
    interface-description [ access-request | accounting-start | accounting-stop ];
    nas-identifier [ access-request | accounting-on | accounting-off | accounting-start |
      accounting-stop ];
    nas-port [ access-request | accounting-start | accounting-stop ];
    nas-port-id [ access-request | accounting-start | accounting-stop ];
    nas-port-type [ access-request | accounting-start | accounting-stop ];
    output-gigapackets [ accounting-stop ];
    output-gigawords [ accounting-stop ];
  }
  ignore {
    framed-ip-netmask;
    input-filter;
    logical-system-routing-instance;
    output-filter;
  }
}
```

Hierarchy Level [edit access profile *profile-name* radius]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify how the router or switch processes RADIUS attributes.

The statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring How RADIUS Attributes Are Used for Subscriber Access

authentication-order

Syntax authentication-order [**ldap radius | none**];

Hierarchy Level [edit **access profile** *profile-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the order of authentication, authorization, and accounting (AAA) servers to use while sending authentication messages.

Default Not enabled

Options **ldap**—Lightweight Directory Access Protocol.

none—No authentication for specified subscribers.

radius—Remote Authentication Dial-In User Service authentication.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch on page 2267
- Configuring 802.1X RADIUS Accounting (CLI Procedure) on page 2339

authentication-order

Syntax authentication-order [*authentication-methods*];

Hierarchy Level [edit access profile *profile-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set the order in which the Junos OS tries different authentication methods when verifying that a client can access the router or switch. For each login attempt, the software tries the authentication methods in order, from first to last.

Default password

Options **password**—Verify the client using the information configured at the **[edit access profile *profile-name* client *client-name*]** hierarchy level.

radius—Verify the client using RADIUS authentication services.

NOTE: For subscriber access management, you must always specify the radius method. Subscriber access management does not support the password keyword (the default), and authentication fails when no method is specified.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Specifying the Authentication and Accounting Methods for Subscriber Access
- Configuring Access Profiles for L2TP or PPP Parameters
- Example: Configuring CHAP Authentication with RADIUS

authentication-profile-name

Syntax authentication-profile-name *access-profile-name*;

Hierarchy Level [edit protocols **dot1x authenticator**],
[edit services **captive-portal**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the name of the access profile to be used for 802.1X, MAC RADIUS, or captive portal authentication.

Default No access profile is specified.

Options ***access-profile-name***—Name of the access profile. The access profile is configured at the [**edit access profile**] hierarchy level and contains the RADIUS server IP address and other information used for authentication.

Required Privilege Level routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation

- Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch on page 2267
- Example: Configuring MAC RADIUS Authentication on a J-EX Series Switch on page 2286
- Example: Setting Up Captive Portal Authentication on a J-EX Series Switch on page 2323
- Configuring 802.1X Interface Settings (CLI Procedure) on page 2331
- Configuring 802.1X Authentication (J-Web Procedure) on page 2332
- Configuring Captive Portal Authentication (CLI Procedure) on page 2350

authentication-server

Syntax authentication-server [*server-addresses*];

Hierarchy Level [edit **access profile** *profile-name* **radius**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the Remote Authentication Dial-In User Service (RADIUS) server for authentication. To configure multiple RADIUS servers, include multiple server addresses. The servers are tried in order and in a round-robin fashion until a valid response is received from one of the servers or until all the configured retry limits are reached.

Default Not enabled

Options ***server-addresses*** —Configure one or more RADIUS server addresses.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch on page 2267
- **show network-access aaa statistics authentication on page 2506**
- Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch on page 2267

authentication-whitelist

Syntax

```
authentication-whitelist {
  mac-address {
  interface interface-name;
  vlan-assignment ( vlan-id |vlan-name);
}
```

Hierarchy Level [edit ethernet-switching-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure MAC addresses for which RADIUS authentication is to be bypassed.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Setting Up Captive Portal Authentication on a J-EX Series Switch on page 2323
- Configuring Captive Portal Authentication (CLI Procedure) on page 2350

authenticator

Syntax

```
authenticator {
  authentication-profile-name access-profile-name;
  interface (all | [ interface-names ]) {
    disable;
    guest-vlan ( vlan-id | vlan-name);
    mac-radius <restrict>;
    maximum-requests number;
    no-reauthentication;
    quiet-period seconds;
    reauthentication {
      interval seconds;
    }
    retries number;
    server-fail (deny | permit | use-cache | vlan-id | vlan-name);
    server-reject-vlan (vlan-id | vlan-name);
    server-timeout seconds;
    supplicant (single | single-secure | multiple);
    supplicant-timeout seconds;
    transmit-period seconds;
  }
  static mac-address {
    interface interface-name;
    vlan-assignment vlan-identifier;
  }
}
```

Hierarchy Level [edit protocols dot1x]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure an authenticator for 802.1X authentication.

The statements are explained separately.

Default No static MAC address or VLAN is configured.

Required Privilege Level routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation

- Configuring 802.1X Interface Settings (CLI Procedure) on page 2331
- Specifying RADIUS Server Connections on a J-EX Series Switch (CLI Procedure) on page 2330
- Example: Configuring Static MAC Bypass of Authentication on a J-EX Series Switch on page 2281
- Understanding Static MAC Bypass of Authentication on J-EX Series Switches

captive-portal

Syntax

```
captive-portal {
  authentication-profile-name authentication-profile-name
  custom-options {
    banner-message string;
    footer-bgcolor color;
    footer-message string;
    form-header-bgcolor color;
    form-header-message string;
    form-reset-label label name;
    form-submit-label label name;
    header-bgcolor color;
    header-logo filename;
    header-message string;
    post-authentication-url url-string;
  }
  interface (all | [interface-names]) {
    quiet-period seconds;
    retries number-of-retries;
    server-timeout seconds;
    reauthentication-timeout seconds;
    supplicant (multiple | single | single-secure);
  }
  secure-authentication (http | https);
}
```

Hierarchy Level [edit services]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure captive portal to authenticate clients connected to the switch for access to the network.

The remaining statements are explained separately.

Default Captive portal is disabled.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Setting Up Captive Portal Authentication on a J-EX Series Switch on page 2323
- Designing a Captive Portal Authentication Login Page on a J-EX Series Switch on page 2351
- Configuring Captive Portal Authentication (CLI Procedure) on page 2350

ca-type

Syntax

```
ca-type {
  number {
    ca-value value;
  }
}
```

Hierarchy Level [edit protocols **lldp-med interface** (all | *interface-name* **location civic-based**)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For Link Layer Discovery Protocol–Media Endpoint Device (LLDP-MED), configure the address elements. These elements are included in the location information to be advertised from the switch to the MED. This information is used during emergency calls to identify the location of the MED.

For further information about the values that can be used to comprise the location,, refer to RFC 4776, *Dynamic Host Configuration Protocol (DHCPv4 and DHCPv6) Option for Civic Addresses Configuration Information*. A subset of those values is provided below.

The **ca-value** statement is explained separately.

Default Disabled.

Options **value**—Civic address elements that represent the civic or postal address. Values are:

- **0**—A code that specifies the language used to describe the location.
- **16**—The leading-street direction, such as “N”.
- **17**—A trailing street suffix, such as “SW”.
- **18**—A street suffix or type, such as “Ave” or “Platz”.
- **19**—A house number, such as “6450”.
- **20**—A house-number suffix, such as “A” or “1/2”.
- **21**—A landmark, such as “Stanford University”.
- **22**—Additional location information, such as “South Wing”.
- **23**—The name and occupant of a location, such as “Carrillo's Holiday Market”.
- **24**—A house-number suffix, such as “95684”.
- **25**—A building structure, such as “East Library”.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show lldp on page 2489**

- Example: Setting Up VoIP with 802.1X and LLDP-MED on a J-EX Series Switch on page 2302
- Configuring LLDP-MED (CLI Procedure) on page 2346

ca-value

Syntax ca-value *value*;

Hierarchy Level [edit protocols **lldp-med interface** (all | *interface-name*) **location civic-based ca-type** *number*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For Link Layer Discovery Protocol–Media Endpoint Device (LLDP-MED), configure location information, such as street address and city, that is indexed by the **ca-type** code. This information is advertised from the switch to the MED and is used during emergency calls to identify the location of the MED.

Default Disabled.

Options *value*—Specify a value that correlates to the **ca-type**. See **ca-type** for a list of codes and suggested values.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show lldp on page 2489**
- Example: Setting Up VoIP with 802.1X and LLDP-MED on a J-EX Series Switch on page 2302
- Configuring LLDP-MED (CLI Procedure) on page 2346

civic-based

Syntax

```
civic-based {
  what number;
  country-code code;
  ca-type {
    number {
      ca-value value;
    }
  }
}
```

Hierarchy Level [edit protocols **lldp-med interface** (all | *interface-name*) **location**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For Link Layer Discovery Protocol–Media Endpoint Discovery (LLDP-MED), configure the geographic location to be advertised from the switch to the MED. This information is used during emergency calls to identify the location of the MED.

The statements are explained separately.

Default Disabled.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show lldp on page 2489**
- Example: Setting Up VoIP with 802.1X and LLDP-MED on a J-EX Series Switch on page 2302
- Configuring LLDP-MED (CLI Procedure) on page 2346

country-code

Syntax country-code *code*;

Hierarchy Level [edit protocols **lldp-med interface** (all | *interface-name*)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For Link Layer Discovery Protocol–Media Endpoint Device (LLDP-MED), configure the two-letter country code to include in the location information. Location information is advertised from the switch to the MED, and is used during emergency calls to identify the location of the MED. The country code is required when configuring LLDP-MED based on location.

Default Disabled.

Options ***code***—Two-letter ISO 3166 country code in capital ASCII letters; for example, US or DE.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show lldp on page 2489**
- Example: Setting Up VoIP with 802.1X and LLDP-MED on a J-EX Series Switch on page 2302
- Configuring LLDP-MED (CLI Procedure) on page 2346

custom-options

Syntax

```
custom-options {
  banner-message string;
  footer-bgcolor color;
  footer-message string;
  form-header-bgcolor color;
  form-header-message string;
  form-reset-label label name;
  form-submit-label label name;
  header-bgcolor color;
  header-logo filename;
  header-message string;
  post-authentication-url url-string;
}
```

Hierarchy Level [edit services **captive-portal**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the design elements of a captive portal login page.

Options

banner-message—The first screen displayed before the captive portal login page is displayed—for example, a disclaimer message.

Range: 1–2047 characters

footer-bgcolor —The hexadecimal color code for the color of the footer bar across the bottom of the captive portal login page—for example, #2E8B57 (sea green).

Values: # symbol followed by six characters.

footer-message—Text message displayed in the footer bar across the bottom of the captive portal login page.

Range: 1–2047 characters

form-header-bgcolor —The hexadecimal color code for the background color of the header bar across the top of the form area of the captive portal login page.

Values: # symbol followed by six characters.

form-header-message—Text message displayed in the header bar across the top of the form area of the captive portal login page.

Range: 1–255 characters

Default: **Captive Portal User Authentication**

form-reset-label—Label displayed in the button that the user can select to clear the username and password fields on the form.

Range: 1–255 characters

Default: **Reset**

form-submit-label —Label displayed in the button that the user selects to submit their login information—for example, **Log In** or **OK**.

Range: 1–255 characters

Default: Log In

header-bgcolor—The hexadecimal color code for the color of the header bar across the top of the captive portal login page.

Values: # symbol followed by six characters.

header-logo—Filename of the file containing the image of the logo displayed at the top of the captive portal login page. The image file can be in GIF, JPEG, or PNG format.

header-message—Text displayed in the header bar across the bottom of the captive portal login page.

Range: 1–2047 characters

Default: User Authentication

post-authentication-url—URL to which the users are directed upon successful authentication—for example **www.mycafe.com**.

Range: 1–255 characters

Default: The page originally requested by the user.

Required Privilege Level

routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation

- Designing a Captive Portal Authentication Login Page on a J-EX Series Switch on page 2351
- Configuring Captive Portal Authentication (CLI Procedure) on page 2350

destination

Syntax

```
destination {
  radius {
    server {
      server-address {
        accounting-port port-number;
        secret password;
        source-address address;
        retry number;
        timeout seconds;
      }
    }
  }
  tacplus {
    server {
      server-address {
        port port-number;
        secret password;
        single-connection;
        timeout seconds;
      }
    }
  }
}
```

Hierarchy Level [edit system accounting]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the authentication server.

Options The remaining statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring RADIUS System Accounting
- Configuring TACACS+ System Accounting

disable

Syntax disable;

Hierarchy Level [edit protocols **dot1x authenticator** interface (all | [*interface-names*])]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable 802.1X authentication on a specified interface or all interfaces.

Default 802.1X authentication is disabled on all interfaces.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show dot1x on page 2477**
- Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on a J-EX Series Switch on page 2290
- Example: Setting Up 802.1X in Conference Rooms to Provide Internet Access to Corporate Visitors on a J-EX Series Switch on page 2276
- Example: Setting Up VoIP with 802.1X and LLDP-MED on a J-EX Series Switch on page 2302
- Example: Configuring Static MAC Bypass of Authentication on a J-EX Series Switch on page 2281
- Configuring 802.1X Interface Settings (CLI Procedure) on page 2331
- Configuring 802.1X Authentication (J-Web Procedure) on page 2332

disable

Syntax disable;

Hierarchy Level [edit protocols lldp],
[edit protocols interface lldp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable the LLDP configuration on the switch or on one or more interfaces.

Default If you do not configure LLDP, it is disabled on the switch and on specific switch interfaces.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show lldp on page 2489**
- Configuring LLDP (CLI Procedure) on page 2344
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261

disable

Syntax disable;

Hierarchy Level [edit protocols lldp-med],
[edit protocols lldp-med interface]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable the LLDP-MED configuration on the switch or on one or more interfaces.

Default If you do not configure LLDP-MED, it is disabled on the switch and on specific switch interfaces.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show lldp on page 2489**
- Configuring LLDP (CLI Procedure) on page 2344
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261

dot1x

Syntax

```
dot1x {
  authenticator {
    authentication-profile-name access-profile-name;
    interface (all | [ interface-names ]) {
      disable;
      guest-vlan (vlan-id | vlan-name);
      mac-radius <restrict>;
      maximum-requests number;
      no-reauthentication;
      quiet-period seconds;
      reauthentication {
        interval seconds;
      }
      retries number;
      server-fail (deny | permit | use-cache | vlan-id | vlan-name);
      server-reject-vlan (vlan-id | vlan-name);
      server-timeout seconds;
      supplicant (single | single-secure | multiple);
      supplicant-timeout seconds;
      transmit-period seconds;
    }
    static mac-address {
      interface interface-names;
      vlan-assignment (vlan-id |vlan-name);
    }
  }
}
```

Hierarchy Level [edit protocols]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure 802.1X authentication for Port-Based Network Access Control. 802.1X authentication is supported on interfaces that are members of private VLANs (PVLANs).

The remaining statements are explained separately.

Default 802.1X is disabled.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show dot1x on page 2477**
- Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on a J-EX Series Switch on page 2290
- Example: Setting Up 802.1X in Conference Rooms to Provide Internet Access to Corporate Visitors on a J-EX Series Switch on page 2276

- Example: Setting Up VoIP with 802.1X and LLDP-MED on a J-EX Series Switch on page 2302
- Example: Configuring Static MAC Bypass of Authentication on a J-EX Series Switch on page 2281
- Example: Configuring MAC RADIUS Authentication on a J-EX Series Switch on page 2286
- Configuring Server Fail Fallback (CLI Procedure) on page 2337

elin

Syntax elin *number*;

Hierarchy Level [edit protocols **lldp-med interface** (all | *interface-name* **location**)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For Link Layer Discovery Protocol–Media Endpoint Discovery (LLDP-MED), configure the Emergency Line Identification Number (ELIN) as location information. Location information is advertised from the switch to the MED device and is used during emergency calls to identify the location of the MED device.

Default Disabled.

Options ***number***—Configure a 10-digit number (area code and telephone number).

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show lldp on page 2489**
- Example: Setting Up VoIP with 802.1X and LLDP-MED on a J-EX Series Switch on page 2302
- Configuring LLDP-MED (CLI Procedure) on page 2346

ethernet-port-type-virtual

Syntax ethernet-port-type-virtual;

Hierarchy Level [edit access profile *profile-name* radius **options**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the physical port type the router or switch uses to authenticate clients. The router or switch passes a port type of **ethernet** in RADIUS attribute 61 (NAS-Port-Type) by default. This statement specifies a port type of **virtual**.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring RADIUS Server Options for Subscriber Access
- Configuring RADIUS Server Parameters for Subscriber Access

ethernet-switching-options

Syntax

```
ethernet-switching-options {
  analyzer {
    name {
      loss-priority priority;
      ratio number;
      input {
        ingress {
          interface (all | interface-name);
          vlan (vlan-id | vlan-name);
        }
        egress {
          interface (all | interface-name);
        }
      output {
        interface interface-name;
        vlan (vlan-id | vlan-name);
      }
    }
  }
  bpdu-block {
    disable-timeout timeout;
    interface (all | [interface-name]);
  }
  dot1q-tunneling {
    ether-type (0x8100 | 0x88a8 | 0x9100);
  }
  interfaces interface-name {
    no-mac-learning;
  }
  mac-notification {
    notification-interval seconds;
  }
  mac-table-aging-time seconds;
  port-error-disable {
    disable-timeout timeout;
  }
  redundant-trunk-group {
    group-name name {
      interface interface-name <primary>;
      interface interface-name;
    }
  }
  secure-access-port {
    dhcp-snooping-file {
      location local_pathname | remote_URL;
      timeout seconds;
      write-interval seconds;
    }
    interface (all | interface-name) {
      allowed-mac {
        mac-address-list;
      }
```

```
        (dhcp-trusted | no-dhcp-trusted);
        mac-limit limit action action;
        no-allowed-mac-log;
        static-ip ip-address {
          vlan  vlan-name;
          mac mac-address;
        }
      }
      vlan (all | vlan-name) {
        (arp-inspection | no-arp-inspection);
        dhcp-option82 {
          circuit-id {
            prefix hostname;
            use-interface-description;
            use-vlan-id;
          }
          remote-id {
            prefix hostname | mac | none;
            use-interface-description;
            use-string string;
          }
          vendor-id [string];
        }
        (examine-dhcp | no-examine-dhcp);
        (ip-source-guard | no-ip-source-guard);
        mac-move-limit limit action action;
      }
    }
    storm-control {
      action-shutdown;
      interface (all | interface-name) {
        bandwidth bandwidth;
        no-broadcast;
        no-unknown-unicast;
      }
    }
    traceoptions {
      file filename <files number> <no-stamp> <replace> <size size> <world-readable |
        no-world-readable>;
      flag flag <disable>;
    }
    unknown-unicast-forwarding {
      vlan (all | vlan-name) {
        interface interface-name;
      }
    }
    voip {
      interface (all | [interface-name | access-ports]) {
        vlan vlan-name ;
        forwarding-class <assured-forwarding | best-effort | expedited-forwarding |
          network-control>;
      }
    }
  }
```

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure Ethernet switching options.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation

- Understanding Port Mirroring on J-EX Series Switches on page 3245
- Port Security for J-EX Series Switches Overview on page 2545
- Understanding BPDU Protection for STP, RSTP, and MSTP on J-EX Series Switches on page 1278
- Understanding Redundant Trunk Links on J-EX Series Switches on page 1049
- Understanding Storm Control on J-EX Series Switches on page 2511
- Understanding 802.1X and VoIP on J-EX Series Switches on page 2263
- Understanding Q-in-Q Tunneling on J-EX Series Switches on page 1051
- Understanding Unknown Unicast Forwarding on J-EX Series Switches on page 2512
- Understanding MAC Notification on J-EX Series Switches on page 1060

events

Syntax events [*events*];

Hierarchy Level [edit system **accounting**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the types of events to track and log.

Options ***events***—Event types; can be one or more of the following:

- **change-log**—Audit configuration changes.
- **interactive-commands**—Audit interactive commands (any command-line input).
- **login**—Audit logins.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring TACACS+ System Accounting

exclude

Syntax

```
exclude {
  accounting-authentic [ accounting-on | accounting-off ];
  accounting-delay-time [ accounting-on | accounting-off ];
  accounting-session-id [ access-request | accounting-on | accounting-off | accounting-stop
    ];
  accounting-terminate-cause [ accounting-off ];
  called-station-id [ access-request | accounting-start | accounting-stop ];
  calling-station-id [ access-request | accounting-start | accounting-stop ];
  class [ accounting-start | accounting-stop ];
  dhcp-gi-address [ access-request | accounting-start | accounting-stop ];
  dhcp-mac-address [ access-request | accounting-start | accounting-stop ];
  output-filter [ accounting-start | accounting-stop ];
  event-timestamp [ accounting-on | accounting-off | accounting-start | accounting-stop
    ];
  framed-ip-address [ accounting-start | accounting-stop ];
  framed-ip-netmask [ accounting-start | accounting-stop ];
  input-filter [ accounting-start | accounting-stop ];
  input-gigapackets [ accounting-stop ];
  input-gigawords [ accounting-stop ];
  interface-description [ access-request | accounting-start | accounting-stop ];
  nas-identifier [ access-request | accounting-on | accounting-off | accounting-start |
    accounting-stop ];
  nas-port [ access-request | accounting-start | accounting-stop ];
  nas-port-id [ access-request | accounting-start | accounting-stop ];
  nas-port-type [ access-request | accounting-start | accounting-stop ];
  output-gigapackets [ accounting-stop ];
  output-gigawords [ accounting-stop ];
}
```

Hierarchy Level [edit access profile *profile-name* radius **attributes**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the router or switch to exclude the specified attributes from the specified type of RADIUS message.

Not all attributes are available in all types of RADIUS messages. By default, the router or switch includes the specified attributes in RADIUS Access-Request, Acct-On, Acct-Off, Acct-Start, and Acct-Stop messages.

Options RADIUS attribute type—RADIUS attribute or Juniper Networks VSA number and name.

- **accounting-authentic**—RADIUS attribute 45, Acct-Authentic.
- **accounting-delay-time**—RADIUS attribute 41, Acct-Delay-Time.
- **accounting-session-id**—RADIUS attribute 44, Acct-Session-Id.
- **accounting-terminate-cause**—RADIUS attribute 49, Acct-Terminate-Cause.
- **called-station-id**—RADIUS attribute 30, Called-Station-Id.
- **calling-station-id**—RADIUS attribute 31, Calling-Station-Id.

- **class**—RADIUS attribute 25, Class.
- **dhcp-gi-address**—Juniper VSA 26-57, DHCP-GI-Address.
- **dhcp-mac-address**—Juniper VSA 26-56, DHCP-MAC-Address.
- **event-timestamp**—RADIUS attribute 55, Event-Timestamp.
- **framed-ip-address**—RADIUS attribute 8, Framed-IP-Address.
- **framed-ip-netmask**—RADIUS attribute 9, Framed-IP-Netmask.
- **input-filter**—Juniper VSA 26-10, Ingress-Policy-Name.
- **input-gigapackets**—Juniper VSA 26-42, Acct-Input-Gigapackets.
- **input-gigawords**—RADIUS attribute 52, Acct-Input-Gigawords.
- **interface-description**—Juniper VSA 26-53, Interface-Desc.
- **nas-identifier**—RADIUS attribute 32, NAS-Identifier.
- **nas-port**—RADIUS attribute 5, NAS-Port.
- **nas-port-id**—RADIUS attribute 87, NAS-Port-Id.
- **nas-port-type**—RADIUS attribute 61, NAS-Port-Type.
- **output-filter**—Juniper VSA 26-11, Egress-Policy-Name.
- **output-gigapackets**—Juniper VSA 25-43, Acct-Output-Gigapackets.
- **output-gigawords**—RADIUS attribute 53, Acct-Output-Gigawords.

RADIUS message type

- **access-request**—RADIUS Access-Accept messages.
- **accounting-off**—RADIUS Accounting-Off messages.
- **accounting-on**—RADIUS Accounting-On messages.
- **accounting-start**—RADIUS Accounting-Start messages.
- **accounting-stop**—RADIUS Accounting-Stop messages.

Required Privilege Level

admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring RADIUS Server Parameters for Subscriber Access

fast-start

Syntax fast-start *count*;

Hierarchy Level [edit protocols lldp-med]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the number of Link Layer Discovery Protocol–Media Endpoint Discovery (LLDP-MED) advertisements sent from the switch in the first second after it has detected an LLDP-MED device (such as an IP telephone).

Options **count**—Number of advertisements.

Range: 1 through 10

Default: 3

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show lldp on page 2489**
- Configuring LLDP-MED (CLI Procedure) on page 2346
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261

forwarding-class

Syntax forwarding-class < assured-forwarding | best-effort | expedited-forwarding |network-control >;

Hierarchy Level [edit **ethernet-switching-options voip interface** <all | *interface-name* | access-ports]>

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For J-EX Series switches, configure the forwarding class used to handle packets on the VoIP interface.

Default Disabled.

Options *class*—Forwarding class:

- **assured-forwarding**— Assured forwarding (AF)—Provides a group of values you can define and includes four subclasses: AF1, AF2, AF3, and AF4, each with three drop probabilities: low, medium, and high.
- **best-effort**—Provides no service profile. For the best effort forwarding class, loss priority is typically not carried in a class-of-service (CoS) value, and random early detection (RED) drop profiles are more aggressive.
- **expedited-forwading**—Provides a low loss, low latency, low jitter, assured bandwidth, end-to-end service.
- **network-control**—Provides a typically high priority because it supports protocol control.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Setting Up VoIP with 802.1X and LLDP-MED on a J-EX Series Switch on page 2302
- Example: Configuring VoIP on a J-EX Series Switch Without Including 802.1X Authentication on page 2309
- Example: Configuring VoIP on a J-EX Series Switch Without Including LLDP-MED Support on page 2315

guest-vlan

Syntax guest-vlan (*vlan-id* | *vlan-name*);

Hierarchy Level [edit protocols **dot1x authenticator interface** (all | [*interface-names*])]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the VLAN to which an interface is moved when no 802.1X supplicants are connected on the interface. The VLAN specified must already exist on the switch.

Default None

Options ***vlan-id***—VLAN tag identifier of the guest VLAN.

vlan-name—Name of the guest VLAN.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Setting Up 802.1X in Conference Rooms to Provide Internet Access to Corporate Visitors on a J-EX Series Switch on page 2276
- Understanding Guest VLANs for 802.1X on J-EX Series Switches on page 2259

hold-multiplier

Syntax hold-multiplier *number*;

Hierarchy Level [edit protocols lldp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the multiplier used in combination with the **advertisement-interval** value to determine the length of time LLDP information is held before it is discarded. The default value is 4 (or 120 seconds).

Default Disabled.

Options *number*—A number used as a multiplier.
Range: 2 through 10
Default: 4 (or 120 seconds)

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show lldp on page 2489**
- Configuring LLDP (CLI Procedure) on page 2344
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261

ignore

Syntax

```
ignore {
  framed-ip-netmask;
  input-filter;
  logical-system-routing-instance;
  output-filter;
}
```

Hierarchy Level [edit access profile *profile-name* radius **attributes**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the router or switch to ignore the specified attributes in RADIUS Access-Accept messages. By default, the router or switch processes the attributes it receives from the external server.

Options

framed-ip-netmask—Framed-IP-Netmask (RADIUS attribute 9).

input-filter—Ingress-Policy-Name (VSA 26-10).

logical-system-routing-instance—Virtual-Router (VSA 26-1).

output-filter—Egress-Policy-Name (VSA 26-11).

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring RADIUS Server Parameters for Subscriber Access

immediate-update

Syntax

```
immediate-update;
```

Hierarchy Level [edit access profile *profile-name* **accounting**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the router or switch to send an Acct-Update message to the RADIUS accounting server on receipt of a response (for example, an ACK or timeout) to the Acct-Start message.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring RADIUS Server Parameters for Subscriber Access
- Configuring How Accounting Statistics Are Collected for Subscriber Access

interface

Syntax

```
interface (all | [ interface-names ]) {
  disable;
  guest-vlan (vlan-name | vlan-id);
  mac-radius <restrict>;
  maximum-requests number;
  no-reauthentication;
  quiet-period seconds;
  reauthentication {
    interval seconds;
  }
  retries number;
  server-fail (deny | permit | use-cache | vlan-id | vlan-name);
  server-reject-vlan (vlan-id | vlan-name);
  server-timeout seconds;
  supplicant (single | single-secure | multiple);
  supplicant-timeout seconds;
  transmit-period seconds;
}
```

Hierarchy Level [edit protocols **dot1x authenticator**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure 802.1X authentication for Port-Based Network Access Control for all interfaces or for specific interfaces.

Options **all**—Configure all interfaces for 802.1X authentication.

[***interface-names***]— List of names of interfaces to configure for 802.1X authentication.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show dot1x on page 2477**
- Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on a J-EX Series Switch on page 2290
- Example: Setting Up 802.1X in Conference Rooms to Provide Internet Access to Corporate Visitors on a J-EX Series Switch on page 2276
- Example: Setting Up VoIP with 802.1X and LLDP-MED on a J-EX Series Switch on page 2302
- Example: Configuring MAC RADIUS Authentication on a J-EX Series Switch on page 2286
- Configuring 802.1X Interface Settings (CLI Procedure) on page 2331
- Configuring 802.1X Authentication (J-Web Procedure) on page 2332

interface-description-format

Syntax interface-description-format (adapter | sub-interface);

Hierarchy Level [edit access profile *profile-name* radius **options**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the information that is included in or omitted from the interface description that the router or switch passes to RADIUS for inclusion in the RADIUS attribute 87 (NAS-Port-Id). By default, the router or switch includes both the subinterface and the adapter in the interface description.

Options **adapter**—Include only the adapter in the interface description.

sub-interface—Include only the subinterface in the interface description.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring RADIUS Server Options for Subscriber Access
- Configuring RADIUS Server Parameters for Subscriber Access

interface (Captive Portal)

Syntax

```
interface (all | [interface-names]) {
  quiet-period seconds;
  reauthentication-timeout seconds;
  retries number-of-retries;
  server-timeout seconds;
  supplicant ( multiple | single | single-secure);
}
```

Hierarchy Level [edit service captive-portal]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure captive portal authentication for all interfaces or for specific interfaces.

Options

all—All interfaces to be configured for captive portal authentication.

[*interface-names*]—List of names of interfaces to be configured for captive portal authentication.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation

- Example: Setting Up Captive Portal Authentication on a J-EX Series Switch on page 2323
- Configuring Captive Portal Authentication (CLI Procedure) on page 2350

interface

Syntax

```
interface (all | interface-name) {
  disable;
}
```

Hierarchy Level [edit protocols lldp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure Link Layer Discovery Protocol (LLDP) on all interfaces or on a specific interface.

Default None

Options

all—All interfaces on the switch.

interface-name—Name of a specific interface.

The remaining statement is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring LLDP (CLI Procedure) on page 2344
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261

interface

Syntax

```
interface (all | interface-name) {
    disable;
    location {
        elin number;
        civic-based {
            what number;
            country-code code;
            ca-type {
                number {
                    ca-value value;
                }
            }
        }
    }
}
```

Hierarchy Level [edit protocols lldp-med]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure Link Layer Discovery Protocol–Media Endpoint Discovery (LLDP-MED) on all interfaces or on a specific interface.

Default Not enabled

Options **all**—All interfaces on the switch.

interface-name—Name of a specific interface.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show lldp on page 2489**
- Example: Setting Up VoIP with 802.1X and LLDP-MED on a J-EX Series Switch on page 2302
- Configuring LLDP-MED (CLI Procedure) on page 2346
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261

interface

Syntax interface [*interface-names*];

Hierarchy Level [edit protocols **dot1x authenticator authentication-profile-name static** *mac-address*], [edit **ethernet-switching-options authentication-whitelist**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure interfaces on which the specified MAC addresses are allowed to bypass RADIUS authentication and allowed to connect to the LAN without authentication.

Options ***interface-names***—List of interfaces.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show dot1x static-mac-address on page 2484**
- Example: Configuring Static MAC Bypass of Authentication on a J-EX Series Switch on page 2281
- Example: Setting Up Captive Portal Authentication on a J-EX Series Switch on page 2323
- Understanding Static MAC Bypass of Authentication on J-EX Series Switches
- Example: Setting Up Captive Portal Authentication on a J-EX Series Switch on page 2323
- Configuring Captive Portal Authentication (CLI Procedure) on page 2350

interface

Syntax

```
interface (all | [interface-name] | access-ports) {
  vlan vlan-name );
  forwarding-class <assured-forwarding | best-effort | expedited-forwarding |
    network-control>;
}
```

Hierarchy Level [edit ethernet-switching-options voip]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable voice over IP (VoIP) for all interfaces or specific interfaces.

Options **all** | ***interface-name*** | **access-ports**—Enable VoIP on all interfaces, on a specific interface, or on all access ports.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Setting Up VoIP with 802.1X and LLDP-MED on a J-EX Series Switch on page 2302
- Example: Configuring VoIP on a J-EX Series Switch Without Including 802.1X Authentication on page 2309
- Example: Configuring VoIP on a J-EX Series Switch Without Including LLDP-MED Support on page 2315

lldp

Syntax

```
lldp {
  disable;
  advertisement-interval seconds;
  fast-start number;
  hold-multiplier number;
  interface (all | [interface-name]) {
    disable;
  }
  lldp-configuration-notification-interval seconds;
  management-address ip-management-address;
  ptopo-configuration-maximum-hold-time seconds;
  ptopo-configuration-trap-interval seconds;
  traceoptions {
    file filename <files number> <size size> <world-readable | no-world-readable> <match
      regex>;
    flag flag (detail | disable | receive | send);
  }
}
```

Hierarchy Level [edit protocols]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure Link Layer Discovery Protocol (LLDP). The switch uses LLDP to advertise its identity and capabilities on a LAN, as well as receive information about other network devices. LLDP is defined in the IEEE standard 802.1AB-2005.

The statements are explained separately.

Default LLDP is enabled.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show lldp on page 2489**
- Configuring LLDP-MED (CLI Procedure) on page 2346
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261

lldp-configuration-notification-interval

Syntax lldp-configuration-notification-interval *seconds*;

Hierarchy Level [edit protocols lldp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify how often SNMP trap notifications are generated as a result of LLDP database changes. If the interval value is 0, trap notifications of database changes are disabled.

Default SNMP trap notifications of LLDP database changes are disabled.

Options ***seconds***—Interval between trap notifications about LLDP database changes.
Range: 0 through 3600

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **show lldp on page 2489**

lldp-med

Syntax

```
lldp-med {
  disable;
  fast-start number;
  interface (all | interface-name) {
    disable;
    location {
      elin number;
      civic-based {
        what number;
        country-code code;
        ca-type {
          number {
            ca-value value;
          }
        }
      }
    }
  }
}
```

Hierarchy Level [edit protocols]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure Link Layer Discovery Protocol–Media Endpoint Discovery. LLDP-MED is an extension of LLDP. The switch uses LLDP-MED to support device discovery of VoIP telephones and to create location databases for these telephone locations for emergency services. LLDP-MED is defined in the standard ANSI/TIA-1057 by the Telecommunications Industry Association (TIA).

The statements are explained separately.

Default Disabled.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show lldp on page 2489**
- Example: Setting Up VoIP with 802.1X and LLDP-MED on a J-EX Series Switch on page 2302
- Configuring LLDP-MED (CLI Procedure) on page 2346

location

Syntax

```
location {
  elin number;
  civic-based {
    what number;
    country-code code;
    ca-type{
      number {
        ca-value value;
      }
    }
  }
}
```

Hierarchy Level [edit protocols **lldp-med interface** (all | *interface-name*)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For Link Layer Discovery Protocol–Media Endpoint Discovery (LLDP-MED), configure the location information. Location information is advertised from the switch to the MED. This information is used during emergency calls to identify the location of the MED.

The statements are explained separately.

Default Disabled.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show lldp on page 2489**
- Example: Setting Up VoIP with 802.1X and LLDP-MED on a J-EX Series Switch on page 2302
- Configuring LLDP-MED (CLI Procedure) on page 2346

mac-radius

Syntax mac-radius <flap-on-disconnect> <restrict>;

Hierarchy Level [edit protocols **dot1x authenticator interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure MAC RADIUS authentication for specific interfaces. MAC RADIUS authentication allows LAN access to permitted MAC addresses. When a new MAC address appears on an interface, the switch consults the RADIUS server to check whether the MAC address is a permitted address. If the MAC address is configured on the RADIUS server, the device is allowed access to the LAN.

If MAC RADIUS is configured, the switch first tries to get a response from the host for 802.1X authentication. If the host is unresponsive, the switch attempts to authenticate using MAC RADIUS.

To restrict authentication to MAC RADIUS only, use the **restrict** option. In restrictive mode, all 802.1X packets are eliminated and the attached device on the interface is considered a nonresponsive host.

Options

flap-on-disconnect—(Optional) When the RADIUS server sends a disconnect message to a supplicant, the switch resets the interface on which the supplicant is authenticated. If the interface is configured for multiple supplicant mode, the switch resets all the supplicants on the specified interface. This option takes effect only when the **restrict** option is also set.

restrict—(Optional) Restricts authentication to MAC RADIUS only. When **mac-radius restrict** is configured the switch drops all 802.1X packets. This option is useful when no other 802.1X authentication methods, such as guest VLAN, are needed on the interface, and eliminates the delay that occurs while the switch determines that a connected device is a non-802.1X-enabled host.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show dot1x on page 2477**
- Example: Configuring MAC RADIUS Authentication on a J-EX Series Switch on page 2286
- Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on a J-EX Series Switch on page 2290
- Configuring MAC RADIUS Authentication (CLI Procedure) on page 2335
- Configuring 802.1X Interface Settings (CLI Procedure) on page 2331
- Understanding MAC RADIUS Authentication on J-EX Series Switches

management-address

Syntax management-address *ip-management-address*;

Hierarchy Level [edit protocols lldp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the management address of the switch to be used in the LLDP Management type, length, and value (TLV) .

Default LLDP Management TLV uses the IP address of the switch's management Ethernet interface (**me0**) or the IP address of the virtual management Ethernet (VME) interface if the switch is a Virtual Chassis.

Options ***ip-management-address***—Management address for the switch.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show lldp on page 2489**
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261
- J-EX Series Switches Interfaces Overview on page 863

maximum-requests

Syntax maximum-requests *number*;

Hierarchy Level [edit protocols dot1x authenticator interface (all | [*interface-names*])]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For 802.1X authentication, configure the maximum number of times an EAPOL request packet is retransmitted to the supplicant before the authentication session times out.

Default Two retransmission attempts

Options ***number***—Number of retransmission attempts.
Range: 1 through 10
Default: 2

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring 802.1X Interface Settings (CLI Procedure) on page 2331
- Configuring 802.1X Authentication (J-Web Procedure) on page 2332

nas-identifier

Syntax nas-identifier *identifier-value*;

Hierarchy Level [edit access profile *profile-name* radius **options**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the value for the client RADIUS attribute 32 (NAS-Identifier). This attribute is used for authentication and accounting requests.

Options ***identifier-value***—String to use for authentication and accounting requests.

Range: 1 to 64 characters

Required Privilege Level admin—To view this statement in the configuration.
admin–control—To add this statement to the configuration.

Related Documentation

- Configuring RADIUS Server Options for Subscriber Access
- Configuring RADIUS Server Parameters for Subscriber Access

nas-port-extended-format

Syntax

```
nas-port-extended-format {
  adapter-width width;
  port-width width;
  slot-width width;
  stacked-vlan-width width;
  vlan-width width;
}
```

Hierarchy Level [edit access profile *profile-name* radius **options**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the RADIUS client to use the extended format for RADIUS attribute 5 (NAS-Port) and specify the width of the fields in the NAS-Port attribute.

Options

adapter-width *width*—Number of bits in the adapter field.

port-width *width*—Number of bits in the port field.

slot-width *width*—Number of bits in the slot field.

stacked-vlan-width *width*—Number of bits in the SVLAN ID field.

vlan-width *width*—Number of bits in the VLAN ID field.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring RADIUS Server Options for Subscriber Access
- Configuring RADIUS Server Parameters for Subscriber Access

no-reauthentication

Syntax no-reauthentication;

Hierarchy Level [edit protocols **dot1x authenticator interface** (all | [*interface-names*])]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For 802.1X authentication, disables reauthentication.

Default Not disabled

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring 802.1X Interface Settings (CLI Procedure) on page 2331
- Configuring 802.1X Authentication (J-Web Procedure) on page 2332
- Understanding Authentication onJ-EX Series Switches on page 2248

options

Syntax

```
options {
  accounting-session-id-format (decimal | description);
  client-accounting-algorithm (direct | round-robin);
  client-authentication-algorithm (direct | round-robin);
  ethernet-port-type-virtual;
  interface-description-format [sub-interface | adapter];
  nas-identifier identifier-value;
  nas-port-extended-format {
    adapter-width width;
    port-width width;
    slot-width width;
    stacked-vlan-width width;
    vlan-width width;
  }
  revert-interval interval;
  vlan-nas-port-stacked-format;
}
```

Hierarchy Level [edit access profile *profile-name* radius]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the options used by RADIUS authentication and accounting servers.

The statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring RADIUS Server Parameters for Subscriber Access
- RADIUS Server Options for Subscriber Access

order

Syntax order [**radius | none**];

Hierarchy Level [edit **access profile** *profile-name* **accounting**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the order of authentication, authorization, and accounting (AAA) servers to use while sending accounting messages and updates.

Default Not enabled

Options **none**—No accounting for specified subscribers.

radius—Remote Authentication Dial-In User Service accounting for specified subscribers.

[radius | none]— Use multiple types of accounting in the order specified. RADIUS accounting is initially used. However, if RADIUS servers are not available, no accounting is done.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch on page 2267
- Configuring 802.1X RADIUS Accounting (CLI Procedure) on page 2339

order

Syntax order [*accounting-method*];

Hierarchy Level [edit access profile *profile-name* **accounting**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set the order in which the Junos OS tries different accounting methods for client activity. When a client logs in, the software tries the accounting methods in the specified order.

Options ***accounting-method***—One or more accounting methods. When a client logs in, the software tries the accounting methods in the following order, from first to last. The only valid value is **radius** for RADIUS accounting.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring Authentication and Accounting Parameters for Subscriber Access

port

Syntax port *port-number*;

Hierarchy Level [edit access **radius-server** *server-address*],
[edit access profile *profile-name* **radius-server** *server-address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the port number on which to contact the RADIUS server.

Options ***port-number***—Port number on which to contact the RADIUS server.
Default: 1812 (as specified in RFC 2865)

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring Router or Switch Interaction with RADIUS Servers
- Configuring Authentication and Accounting Parameters for Subscriber Access

port (RADIUS Server)

Syntax port *port-number*;

Hierarchy Level [edit system radius-server *address*],
[edit system accounting destination **radius server** *address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the port number on which to contact the RADIUS server.

Options ***number***—Port number on which to contact the RADIUS server.
Default: **1812** (as specified in RFC 2865)

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring RADIUS Authentication

port (TACACS+ Server)

Syntax port *port-number*;

Hierarchy Level [edit system accounting destination tacplus server *server-address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the port number on which to contact the TACACS+ server.

Options *number*—Port number on which to contact the TACACS+ server.
Default: 49

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring TACACS+ System Accounting

profile

Syntax

```
profile profile-name {
  accounting {
    order [radius | none];
    accounting-stop-on-access-deny;
    accounting-stop-on-failure;
  }
  authentication-order [authentication-method];
  radius {
    accounting-server [server-addresses];
    authentication-server [server-addresses];
  }
}
```

Hierarchy Level [edit access]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure an access profile. The access profile contains the entire authentication, authorization, and accounting (AAA) configuration that aids in handling AAA requests, including the authentication method and order, AAA server addresses, and AAA accounting.

Default Not enabled

Options *profile-name* —Profile name of up to 32 characters.

The remaining statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch on page 2267
- Configuring 802.1X RADIUS Accounting (CLI Procedure) on page 2339

ptopo-configuration-maximum-hold-time

Syntax ptopo-configuration-maximum-hold-time *seconds*;

Hierarchy Level [edit protocols lldp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure how long to maintain the physical topology database entries. The physical topology identifies the devices on the network and their physical interconnections.

Options ***seconds***—Time to maintain physical topology database entries.
Default: 300
Range: 1 through 2147483647

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **show lldp on page 2489**
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261

ptopo-configuration-trap-interval

Syntax ptopo-configuration-trap-interval *seconds*;

Hierarchy Level [edit protocols lldp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify how often SNMP trap notifications are sent regarding changes in physical topology global statistics.

Default SNMP trap notifications of changes in physical topology global statistics are disabled.

Options ***seconds***—Interval between SNMP trap notifications about physical topology global statistics.
Range: 0 through 3600

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

quiet-period

Syntax quiet-period *seconds*;

Hierarchy Level [edit protocols **dot1x authenticator interface** (all | [*interface-names*])]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For 802.1X authentication, configure the number of seconds the interface remains in the wait state following a failed authentication attempt by a supplicant before reattempting authentication.

Default 60 seconds

Options ***seconds***—Number of seconds the interface remains in the wait state.
Range: 0 through 65,535 seconds
Default: 60 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **show network-access aaa statistics authentication on page 2506**
- Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch on page 2267

quiet-period (Captive Portal)

Syntax quiet-period *seconds*;

Hierarchy Level [edit services **captive-portal interface** (all | *interface-names*)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure time, in seconds, after a user exceeds the maximum number of retries before they can attempt to authenticate.

Options ***seconds***—Number of seconds.
Range: 1–65535
Default: 60

Required Privilege Level routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation
- Example: Setting Up Captive Portal Authentication on a J-EX Series Switch on page 2323
- Configuring Captive Portal Authentication (CLI Procedure) on page 2350

radius

Syntax

```
radius {
  accounting-server [server-addresses];
  authentication-server [server-addresses];
}
```

Hierarchy Level [edit **access profile** *profile-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the Remote Authentication Dial-In User Service (RADIUS) servers for authentication and for accounting. To configure multiple RADIUS servers, include multiple **radius**statements. The servers are tried in order and in a round-robin fashion until a valid response is received from one of the servers or until all the configured retry limits are reached.

The statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch on page 2267
- Configuring 802.1X RADIUS Accounting (CLI Procedure) on page 2339
- Filtering 802.1X Supplicants Using RADIUS Server Attributes on page 2340
- Configuring 802.1X RADIUS Accounting (CLI Procedure) on page 2339

radius (Access Profile)

Syntax

```
radius {
    accounting-server [ ip-address ];
    attributes {
        exclude
            accounting-authentic [ accounting-on | accounting-off ];
            accounting-delay-time [ accounting-on | accounting-off ];
            accounting-session-id [ access-request | accounting-on | accounting-off |
                accounting-stop ];
            accounting-terminate-cause [ accounting-off ];
            called-station-id [ access-request | accounting-start | accounting-stop ];
            calling-station-id [ access-request | accounting-start | accounting-stop ];
            class [ accounting-start | accounting-stop ];
            dhcp-gi-address [ access-request | accounting-start | accounting-stop ];
            dhcp-mac-address [ access-request | accounting-start | accounting-stop ];
            output-filter [ accounting-start | accounting-stop ];
            event-timestamp [ accounting-on | accounting-off | accounting-start | accounting-stop
                ];
            framed-ip-address [ accounting-start | accounting-stop ];
            framed-ip-netmask [ accounting-start | accounting-stop ];
            input-filter [ accounting-start | accounting-stop ];
            input-gigapackets [ accounting-stop ];
            input-gigawords [ accounting-stop ];
            interface-description [ access-request | accounting-start | accounting-stop ];
            nas-identifier [ access-request | accounting-on | accounting-off | accounting-start |
                accounting-stop ];
            nas-port [ access-request | accounting-start | accounting-stop ];
            nas-port-id [ access-request | accounting-start | accounting-stop ];
            nas-port-type [ access-request | accounting-start | accounting-stop ];
            output-gigapackets [ accounting-stop ];
            output-gigawords [ accounting-stop ];
        }
        ignore {
            framed-ip-netmask;
            input-filter;
            logical-system-routing-instance;
            output-filter;
        }
    }
    authentication-server [ ip-address ];
    options {
        accounting-session-id-format (decimal | description);
        client-accounting-algorithm (direct | round-robin);
        client-authentication-algorithm (direct | round-robin);
        ethernet-port-type-virtual;
        interface-description-format [sub-interface | adapter];
        nas-identifier identifier-value;
        nas-port-extended-format {
            adapter-width width;
            port-width width;
            slot-width width;
            stacked-vlan-width width;
            vlan-width width;
```

```
    }
    revert-interval interval;
    vlan-nas-port-stacked-format;
  }
}
```

Hierarchy Level [edit access profile *profile-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the RADIUS parameters that the router uses for AAA authentication and accounting for subscribers.

The statements are explained separately.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring RADIUS Server Parameters for Subscriber Access
- RADIUS Server Options for Subscriber Access

radius

Syntax

```
radius {
  server {
    server-address {
      accounting-port port-number;
      secret password;
      source-address address;
      retry number;
      timeout seconds;
    }
  }
}
```

Hierarchy Level [edit system accounting destination]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the RADIUS accounting server.

Options ***server-address***—Address of the RADIUS accounting server.

The remaining statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring RADIUS System Accounting

radius-server

Syntax

```
radius-server server-address {
    accounting-port port-number;
    port port-number;
    retry attempts;
    routing-instance routing-instance-name;
    secret password;
    source-address source-address;
    timeout seconds;
}
```

Hierarchy Level

[edit access],
[edit access profile *profile-name*]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure RADIUS for subscriber access management, L2TP, or PPP.

To configure multiple RADIUS servers, include multiple **radius-server** statements. The servers are tried in order and in a round-robin fashion until a valid response is received from one of the servers or until all the configured retry limits are reached.

Options

server-address—Address of the RADIUS authentication server.

The remaining statements are explained separately.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring RADIUS Authentication for L2TP
- Configuring the PPP Authentication Protocol
- Configuring RADIUS Authentication
- Configuring Authentication and Accounting Parameters for Subscriber Access

reauthentication

Syntax

```
reauthentication {
    interval seconds;
}
```

Hierarchy Level [edit protocols **dot1x authenticator interface** (all | [*interface-names*])]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For 802.1X authentication, specify reauthentication parameters.

Default 3600 seconds.

Options **disable**—Disables the periodic reauthentication of the supplicant.

interval *seconds* —Sets the periodic reauthentication time interval. The range is 1 through 65,535 seconds.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring 802.1X Interface Settings (CLI Procedure) on page 2331
- Configuring 802.1X Authentication (J-Web Procedure) on page 2332
- Understanding Authentication on J-EX Series Switches on page 2248

retries

Syntax retries *number*;

Hierarchy Level [edit protocols **dot1x authenticator interface** (all | [*interface-names*])]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For 802.1X authentication, configure the number of times the switch attempts to authenticate the port after an initial failure. The port remains in a wait state during the quiet period after the authentication attempt.

Default 3 retries

Options ***number*** —Number of retries.
Range: 1 through 10
Default: 3

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring 802.1X Interface Settings (CLI Procedure) on page 2331
- Configuring 802.1X Authentication (J-Web Procedure) on page 2332
- Understanding Authentication on J-EX Series Switches on page 2248

retries (Captive Portal)

Syntax retries *number-of-tries*;

Hierarchy Level [edit services **captive-portal interface** (all | *interface-names*)]]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the number of times the user can attempt to submit authentication information.

Options ***number-of-tries***—Number of authentication attempts by user.
Range: 1–65535
Default: 3

Required Privilege Level routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation
- Example: Setting Up Captive Portal Authentication on a J-EX Series Switch on page 2323
- Configuring Captive Portal Authentication (CLI Procedure) on page 2350

retry

Syntax retry *attempts*;

Hierarchy Level [edit access **radius-server** *server-address*],
[edit access profile *profile-name* **radius-server** *server-address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the number of times that the router or switch is allowed to attempt to contact a RADIUS authentication or accounting server.

Options ***attempts***—Number of times that the router is allowed to attempt to contact a RADIUS server.

Range: 1 through 10

Default: 3

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Authentication and Accounting Parameters for Subscriber Access
- Configuring Router or Switch Interaction with RADIUS Servers
- Example: Configuring CHAP Authentication with RADIUS
- Configuring RADIUS Authentication for L2TP
- **timeout on page 2454**

retry

Syntax retry *number*;

Hierarchy Level [edit system radius-server *server-address*],
[edit system accounting destination radius server *server-address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Number of times the router or switch is allowed to try to contact a RADIUS authentication or accounting server.

Options *number*—Number of retries allowed for contacting a RADIUS server.

Range: 1 through 10

Default: 3

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring RADIUS Authentication
- Configuring RADIUS System Accounting
- **timeout on page 2453**

revert-interval

Syntax revert-interval *interval*;

Hierarchy Level [edit access profile *profile-name* radius **options**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the amount of time the router or switch waits after a server has become unreachable. The router or switch rechecks the connection to the server when the specified interval expires. If the server is then reachable, it is used in accordance with the order of the server list.

Options *interval*—Amount of time to wait.

Range: 0 through 4294967295 seconds

Default: 60 seconds

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring RADIUS Server Options for Subscriber Access
- Configuring Authentication and Accounting Parameters for Subscriber Access

routing-instance

Syntax routing-instance *routing-instance-name*;

Hierarchy Level [edit access **radius-server** *server-address*],
[edit access profile *profile-name* **radius-server** *server-address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the routing instance used to send RADIUS packets to the RADIUS server.

Options ***routing-instance-name***—Routing instance name.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring the PPP Authentication Protocol
- Configuring Authentication and Accounting Parameters for Subscriber Access

secret

Syntax secret *password*;

Hierarchy Level [edit access profile *profile-name* **radius-server** *server-address*],
[edit access radius-disconnect *client-address*],
[edit access **radius-server** *server-address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the password to use with the RADIUS server. The secret password used by the local router or switch must match that used by the server.

Options ***password***—Password to use; it can include spaces if the character string is enclosed in quotation marks.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring Authentication and Accounting Parameters for Subscriber Access
- Configuring Router or Switch Interaction with RADIUS Servers
- Example: Configuring CHAP Authentication with RADIUS
- Configuring RADIUS Authentication for L2TP
- Configuring the RADIUS Disconnect Server for L2TP

secret

Syntax secret *password*;

Hierarchy Level [edit system accounting destination radius **server** *server-address*],
[edit system accounting destination tacplus **server** *server-address*],
[edit system radius-server *server-address*],
[edit system **tacplus-server** *server-address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the password to use with the RADIUS or TACACS+ server. The secret password used by the local router or switch must match that used by the server.

Options ***password***—Password to use; can include spaces included in quotation marks.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring RADIUS Authentication
- Configuring TACACS+ Authentication
- Configuring TACACS+ System Accounting
- Configuring RADIUS System Accounting

secure-authentication

Syntax secure-authentication (http | https);

Hierarchy Level [edit services **captive-portal**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable HTTP or HTTPS access on the captive portal interface.

Default **http**

Options **http**—Enables HTTP access on the captive portal interface.

https—Enables HTTPS access on the captive portal interface. HTTPS is recommended.

Required Privilege Level routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation
- Example: Setting Up Captive Portal Authentication on a J-EX Series Switch on page 2323
- Configuring Captive Portal Authentication (CLI Procedure) on page 2350

server (RADIUS Accounting)

Syntax

```
server {
  server-address {
    accounting-port port-number;
    retry number
    secret password;
    source-address address;
    timeout seconds;
  }
}
```

Hierarchy Level [edit system accounting destination radius]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure RADIUS logging.

The remaining statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring RADIUS System Accounting

server (TACACS+ Accounting)

Syntax

```
server {
  server-address {
    port port-number;
    secret password;
    single-connection;
    timeout seconds;
  }
}
```

Hierarchy Level [edit system accounting destination tacplus]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure TACACS+ logging.

The remaining statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring TACACS+ System Accounting

server-fail

Syntax server-fail (deny | permit | use-cache | *vlan-id* | *vlan-name*);

Hierarchy Level [edit protocols **dot1x authenticator interface** (all | [*interface-names*])]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For J-EX Series switches configured for 802.1X authentication, specify the server fail fallback action the switch takes when all RADIUS authentication servers are unreachable.

When you specify the action ***vlan-name*** or ***vlan-id***, the VLAN must already be configured on the switch.

Default Authentication is denied.

Options **deny**—Force fail the supplicant authentication. No traffic will flow through the interface.

permit—Force succeed the supplicant authentication. Traffic will flow through the interface as if it were successfully authenticated by the RADIUS server.

use-cache—Force succeed the supplicant authentication only if it was previously authenticated successfully. This action ensures that already authenticated supplicants are not affected.

vlan-id—Move supplicant on the interface to the VLAN specified by this numeric identifier. This action is allowed only if it is the first supplicant connecting to the interface. If an authenticated supplicant is already connected, then the supplicant is not moved to the VLAN and is not authenticated.

vlan-name—Move supplicant on the interface to the VLAN specified by this name. This action is allowed only if it is the first supplicant connecting to an interface. If an authenticated supplicant is already connected, then the supplicant is not moved to the VLAN and is not authenticated.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show dot1x on page 2477**
- Example: Configuring 802.1X Authentication Options When the RADIUS Server is Unavailable to a J-EX Series Switch on page 2271
- Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch on page 2267
- Configuring Server Fail Fallback (CLI Procedure) on page 2337
- Understanding Server Fail Fallback and 802.1X Authentication on J-EX Series Switches on page 2258

server-reject-vlan

Syntax server-reject-vlan (*vlan-id* | *vlan-name*);

Hierarchy Level [edit protocols **dot1x authenticator interface** (all | [*interface-names*])]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For J-EX Series switches configured for 802.1X authentication, specify that when the switch receives an Extensible Authentication Protocol Over LAN (EAPOL) Access-Reject message during the authentication process between the switch and the RADIUS authentication server, supplicants attempting access to the LAN are granted access and moved to a specific VLAN. Any VLAN name or VLAN ID sent by a RADIUS server as part of the EAPOL Access-Reject message is ignored.

When you specify the VLAN ID or VLAN name, the VLAN must already be configured on the switch.

Default None

Options ***vlan-id*** —Numeric identifier of the VLAN to which the supplicant is moved.

vlan-name—Name of the VLAN to which the supplicant is moved.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show dot1x on page 2477**
- Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch on page 2267
- Configuring Server Fail Fallback (CLI Procedure) on page 2337
- Understanding Server Fail Fallback and 802.1X Authentication on J-EX Series Switches on page 2258

server-timeout

Syntax server-timeout *seconds*;

Hierarchy Level [edit protocols **dot1x authenticator interface** (all | [*interface-name*])

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For 802.1X authentication, configure the amount of time a port will wait for a reply when relaying a response from the supplicant to the authentication server before timing out and invoking the server-fail action.

Default 30 seconds

Options ***seconds*** —Number of seconds.
Range: 1 through 60 seconds
Default: 30 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show dot1x on page 2477**
- **clear dot1x on page 2468**
- Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch on page 2267
- 802.1X for J-EX Series Switches Overview on page 2253

server-timeout (Captive Portal)

Syntax server-timeout *seconds*;

Hierarchy Level [edit services **captive-portal interface** (all | *interface-names*)]]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the time in seconds an interface will wait for a reply when relaying a response from the client to the authentication server before timing out and invoking the server-fail action.

Options ***seconds***—Number of seconds.
Range: 1–65535
Default: 20

Required Privilege Level routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation
- Example: Setting Up Captive Portal Authentication on a J-EX Series Switch on page 2323
- Configuring Captive Portal Authentication (CLI Procedure) on page 2350

session-expiry

Syntax session-expiry *seconds*;

Hierarchy Level [edit services **captive-portal interface** (all | *interface-names*)]]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the maximum duration in seconds of a session.

Options ***seconds***—Duration of session.
Range: 1 through 65535
Default: 3600

Required Privilege Level routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation
- Example: Setting Up Captive Portal Authentication on a J-EX Series Switch on page 2323
- Configuring Captive Portal Authentication (CLI Procedure) on page 2350

single-connection

Syntax single-connection;

Hierarchy Level [edit system accounting destination tacplus-server *server-address*]
[edit system **tacplus-server** *server-address*],

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Optimize attempts to connect to a TACACS+ server. The software maintains one open TCP connection to the server for multiple requests rather than opening a connection for each connection attempt.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Configuring TACACS+ Authentication
- Configuring TACACS+ System Accounting

source-address

Syntax source-address *source-address*;

Hierarchy Level [edit access **radius-server** *server-address*],
[edit access profile *profile-name* **radius-server** *server-address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a source address for each configured RADIUS server. Each RADIUS request sent to a RADIUS server uses the specified source address.

Options ***source-address***—A valid IPv4 address configured on one of the router or switch interfaces. On M Series routers only, the source address can be an IPv6 address and the UDP source port is 514.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation
- Configuring Router or Switch Interaction with RADIUS Servers
- Configuring Authentication and Accounting Parameters for Subscriber Access
- Example: Configuring CHAP Authentication with RADIUS
- Configuring RADIUS Authentication for L2TP

source-address (NTP, RADIUS, System Logging, or TACACS+)

Syntax source-address *source-address*;

Hierarchy Level [edit system accounting destination radius **server** *server-address*],
[edit system accounting destination tacplus **server** *server-address*],
[edit system **ntp**],
[edit system radius-server *server-address*],
[edit system **syslog**],
[edit system **tacplus-server** *server-address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify a source address for each configured TACACS+ server, RADIUS server, NTP server, or the source address to record in system log messages that are directed to a remote machine.

Options ***source-address***—A valid IP address configured on one of the router or switch interfaces. For system logging, the address is recorded as the message source in messages sent to the remote machines specified in all **host *hostname*** statements at the **[edit system syslog]** hierarchy level, but not for messages directed to the other Routing Engine..

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring RADIUS Authentication
- Synchronizing and Coordinating Time Distribution Using NTP
- Specifying an Alternative Source Address for System Log Messages

static

Syntax

```
static mac-address {
  interface interface-names;
  vlan-assignment (vlan-id |vlan-name );
}
```

Hierarchy Level [edit protocols **dot1x authenticator authentication-profile-name**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure MAC addresses to exclude from 802.1X authentication. The static MAC list provides an authentication bypass mechanism for supplicants connecting to a port, permitting devices such as printers that are not 802.1X-enabled to be connected to the network on 802.1X-enabled ports.

Using this 802.1X authentication-bypass mechanism, the supplicant connected to the MAC address is assumed to be successfully authenticated and the port is opened for it. No further authentication is done for the supplicant.

You can optionally configure the VLAN that the supplicant is moved to or the interfaces on which the MAC address can gain access from.

Options ***mac-address*** —The MAC address of the device for which 802.1X authentication should be bypassed and the device permitted access to the port.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show dot1x static-mac-address on page 2484**
- Example: Configuring Static MAC Bypass of Authentication on a J-EX Series Switch on page 2281
- Configuring 802.1X Interface Settings (CLI Procedure) on page 2331
- Configuring 802.1X Authentication (J-Web Procedure) on page 2332
- Understanding Static MAC Bypass of Authentication on J-EX Series Switches

statistics

Syntax statistics (time | volume-time);

Hierarchy Level [edit access profile *profile-name* accounting]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the router or switch to collect time statistics, or both volume and time statistics, for the sessions being managed by AAA.

Options **time**—Collect uptime statistics only.

volume-time—Collect both volume and uptime statistics. This option is not available for Mobile IP.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Mobile IP Home Agent Elements and Behavior
- Configuring Authentication and Accounting Parameters for Subscriber Access

supplicant

Syntax supplicant (multiple | single | single-secure);

Hierarchy Level [edit protocols **dot1x authenticator interface** (all | [*interface-names*])],
[edit services **captive-portal interface** (all | *interface-names*)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the method used to authenticate clients for 802.1X or captive portal authentication.

Default **single**

Options

single—Authenticates only the first client that connects to an authenticator port. All other clients connecting to the authenticator port after the first are permitted free access to the port without further authentication. If the first authenticated client logs out, all other supplicants are locked out until a client authenticates again.

single-secure—Authenticates only one client to connect to an authenticator port. The host must be directly connected to the switch.

multiple—Authenticates multiple clients individually on one authenticator port. You can configure the number of clients per port. If you also configure a maximum number of devices that can be connected to a port through port security settings, the lower of the configured values is used to determine the maximum number of clients allowed per port.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on a J-EX Series Switch on page 2290
- Example: Setting Up Captive Portal Authentication on a J-EX Series Switch on page 2323
- Understanding Authentication on J-EX Series Switches on page 2248
- Understanding Captive Portal Authentication
- Configuring Captive Portal Authentication (CLI Procedure) on page 2350

supplicant-timeout

Syntax supplicant-timeout *seconds*;

Hierarchy Level [edit protocols **dot1x authenticator interface** (all | [*interface-name*])

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For 802.1X authentication, configure how long the port waits for a response when relaying a request from the authentication server to the supplicant before resending the request.

Default 30 seconds

Options ***seconds*** —Number of seconds.
Range: 1 through 60 seconds
Default: 30 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **supplicant on page 2450**
- Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on a J-EX Series Switch on page 2290
- Understanding Authentication on J-EX Series Switches on page 2248

tacplus

Syntax

```
tacplus {
  server {
    server-address {
      port port-number;
      secret password;
      single-connection;
      timeout seconds;
    }
  }
}
```

Hierarchy Level [edit system accounting destination]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the Terminal Access Controller Access Control System Plus (TACACS+).

Options ***server-address***—Address of the TACACS+ authentication server.

The remaining statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring TACACS+ System Accounting

timeout

Syntax timeout *seconds*;

Hierarchy Level [edit system radius-server *server-address*],
[edit system **tacplus-server** *server-address*],
[edit system accounting destination radius server server-address],
[edit system accounting destination tacplus server server-address]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the amount of time that the local router or switch waits to receive a response from a RADIUS or TACACS+ server.

Options ***seconds***—Amount of time to wait.

Range: 1 through 90 seconds

Default: 3 seconds

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring RADIUS Authentication
- Configuring TACACS+ Authentication
- **retry on page 2438**

timeout (RADIUS)

Syntax timeout *seconds*;

Hierarchy Level [edit access **radius-server** *server-address*],
[edit access profile *profile-name* **radius-server** *server-address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the amount of time that the local router or switch waits to receive a response from a RADIUS server.

Options ***seconds***—Amount of time to wait.

Range: 1 through 90 seconds

Default: 3 seconds

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Configuring Router or Switch Interaction with RADIUS Servers
- Configuring Authentication and Accounting Parameters for Subscriber Access
- Example: Configuring CHAP Authentication with RADIUS
- Configuring RADIUS Authentication for L2TP

traceoptions

Syntax

```
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable> <match
    regex>;
  flag flag ;
}
```

Hierarchy Level [edit protocols dot1x]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define tracing operations for the 802.1X protocol.

Default Tracing operations are disabled.

Options

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

file *number*—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify gigabytes number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **size**option.

Range: 2 through 1000

Default: 3 files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:

- **all**—All tracing operations.
- **config-internal**—Trace internal configuration operations.
- **general**—Trace general operations.
- **normal**—Trace normal operations.
- **parse**—Trace reading of the configuration.
- **regex-parse**—Trace regular-expression parsing operations.
- **state**—Trace protocol state changes.
- **task**—Trace protocol task operations.
- **timer**—Trace protocol timer operations.

match *regex*—(Optional) Refine the output to include lines that contain the regular expression.

no-world-readable—(Optional) Restricted file access to the user who created the file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **files** option.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify gigabyte

Range: 10 KB through 1gigabyte

Default: 128 KB

world-readable—(Optional) Enable unrestricted file access.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show lldp on page 2489**
- Configuring 802.1X Interface Settings (CLI Procedure) on page 2331
- 802.1X for J-EX Series Switches Overview on page 2253

traceoptions

Syntax

```
traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable> <match
    regex>;
  flag flag (detail | disable | receive | send);
}
```

Hierarchy Level [edit protocols lldp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define tracing operations for the LLDP protocol.

Default Tracing operations are disabled.

Options

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000

Default: 3 files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:

- **all**—All tracing operations.
- **config**—Trace configuration operations.
- **packet**—Trace packet events.
- **rtsock**—Trace routing socket operations.

match *regex*—(Optional) Refine the output to include lines that contain the regular expression.

no-world-readable—(Optional) Restrict file access to the user who created the file.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **files** option.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10 KB through 1 GB
Default: 128 KB

world-readable—(Optional) Enable unrestricted file access.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring LLDP-MED (CLI Procedure) on page 2346
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261

transmit-delay

Syntax transmit-delay *seconds*;

Hierarchy Level [edit protocols lldp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the delay between 2 successive LLDP advertisements.

Default Disabled.

Options *seconds*—Number of seconds between two successive LLDP advertisements.
Range: 1 through 8192 seconds
Default: 2

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- **show lldp on page 2489**
- Configuring LLDP (CLI Procedure) on page 2344
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261

transmit-period

Syntax transmit-period *seconds*;

Hierarchy Level [edit protocols **dot1x authenticator interface** (all | [*interface-name*])

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For 802.1X authentication, how long the port waits before retransmitting the initial EAPOL PDUs to the supplicant.

Default 30 seconds

Options *seconds*—Number of seconds the port waits before retransmitting the initial EAPOL PDUs to the supplicant.

Range: 1 through 65,535 seconds

Default: 30 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring 802.1X Interface Settings (CLI Procedure) on page 2331
- 802.1X for J-EX Series Switches Overview on page 2253

update-interval

Syntax update-interval *minutes*;

Hierarchy Level [edit access profile *profile-name* **accounting**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the amount of time that the router or switch waits before sending a new accounting update.

Options ***minutes***—Amount of time between updates, in minutes.

Range: 10 through 1440 minutes

Default: No updates

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring Authentication and Accounting Parameters for Subscriber Access

vlan-assignment

Syntax vlan-assignment (*vlan-id* | *vlan-name*);

Hierarchy Level [edit protocols **dot1x authenticator authentication-profile-name static** *mac-address*], [edit **ethernet-switching-options authentication-whitelist**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the VLAN that is associated with the list of MAC addresses that are excluded from RADIUS authentication.

Options ***vlan-id*** | ***vlan-name***—The name of the VLAN or the VLAN tag identifier to associate with the device. The VLAN already exists on the switch.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show dot1x static-mac-address on page 2484**
- Example: Configuring Static MAC Bypass of Authentication on a J-EX Series Switch on page 2281
- Example: Setting Up Captive Portal Authentication on a J-EX Series Switch on page 2323
- Understanding Static MAC Bypass of Authentication on J-EX Series Switches
- Example: Setting Up Captive Portal Authentication on a J-EX Series Switch on page 2323
- Configuring Captive Portal Authentication (CLI Procedure) on page 2350

vlan-nas-port-stacked-format

Syntax vlan-nas-port-stacked-format;

Hierarchy Level [edit access profile *profile-name* radius **options**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure RADIUS attribute 5 (NAS-Port) to include the S-VLAN ID, in addition to the VLAN ID, for subscribers on Ethernet interfaces.

Required Privilege Level admin—To view this statement in the configuration.
admin-control—To add this statement to the configuration.

Related Documentation

- Configuring RADIUS Server Options for Subscriber Access
- Configuring Authentication and Accounting Parameters for Subscriber Access

vlan

Syntax vlan (*vlan-id* | *vlan-name* | untagged);

Hierarchy Level [edit **ethernet-switching-options voip interface** (all | [*interface-name* | access-ports])

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For J-EX Series switches, specify the VLAN name or VLAN tag identifier associated with the VLAN to be sent from the authenticating server to the IP phone.

Options *vlan-name*—Name of a VLAN.

vlan-id—The VLAN tag identifier.

Range: 0 through 4095. Tags 0 and 4095 are reserved by the Junos OS, and you should not configure them.

untagged—Allow untagged VLAN traffic.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Setting Up VoIP with 802.1X and LLDP-MED on a J-EX Series Switch on page 2302
- Example: Configuring VoIP on a J-EX Series Switch Without Including 802.1X Authentication on page 2309
- Example: Configuring VoIP on a J-EX Series Switch Without Including LLDP-MED Support on page 2315

voip

Syntax

```
voip {
  interface (all | [interface-name | access-ports]) {
    vlan vlan-name );
    forwarding-class <assured-forwarding | best-effort | expedited-forwarding |
      network-control>;
  }
}
```

Hierarchy Level [edit ethernet-switching-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure voice over IP (VoIP) interfaces.

The statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Setting Up VoIP with 802.1X and LLDP-MED on a J-EX Series Switch on page 2302
- Example: Configuring VoIP on a J-EX Series Switch Without Including 802.1X Authentication on page 2309
- Example: Configuring VoIP on a J-EX Series Switch Without Including LLDP-MED Support on page 2315

what

Syntax what *number*;

Hierarchy Level [edit protocols **lldp-med interface** (all | *interface-name*) **location civic-based**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description For Link Layer Discovery Protocol–Media Endpoint Device (LLDP-MED), configure the location to which the DHCP entry refers. This information is advertised, along with other location information, from the switch to the MED. It is used during emergency calls to identify the location of the MED.

Options **0** and **1** should not be used unless it is known that the DHCP client is in close physical proximity to the server or network element.

Default 1

Options *number*—Location:

- **0**—Location of the DHCP server.
- **1**—Location of a network element believed to be closest to the client.
- **2**—Location of the client.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show lldp on page 2489**
- Example: Setting Up VoIP with 802.1X and LLDP-MED on a J-EX Series Switch on page 2302
- Configuring LLDP-MED (CLI Procedure) on page 2346

CHAPTER 86

Operational Commands for 802.1X

clear captive-portal

Syntax **clear captive-portal (firewall** [*interface-names*] **| interface (all |** [*interface-names*]**) | mac-address** [*mac-addresses*])

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Reset the authentication state of a captive portal interface or captive-portal firewall statistics on one or more interfaces.

Options

firewall [*interface-names*]—Resets captive portal statistics on all interfaces or on the specified interface.

interface (all | *interface-names*)—Resets the authentication state of users connected to all interfaces or the specified interfaces.

mac-address *mac-addresses*—Resets the authentication state for the specified MAC addresses.

Required Privilege Level view

Related Documentation

- **show captive-portal authentication-failed-users on page 2471**
- **show captive-portal interface on page 2474**
- **show captive-portal firewall on page 2472**
- Example: Setting Up Captive Portal Authentication on a J-EX Series Switch on page 2323
- Configuring Captive Portal Authentication (CLI Procedure) on page 2350

List of Sample Output

clear captive-portal interface on page 2467
clear captive-portal interface on page 2467
clear captive-portal mac-address on page 2467
clear captive-portal firewall on page 2467

Output Fields Table 304 on page 2466 lists the output fields for the **clear captive-portal interface** command. (The **clear captive-portal firewall** and **clear captive-portal mac-address** commands have no output). Output fields are listed in the approximate order in which they appear.

Table 304: clear captive-portal interface Output Fields

Field Name	Field Description
Interface	Interface on which captive portal has been configured.

Table 304: clear captive-portal interface Output Fields *(continued)*

Field Name	Field Description
State	The state of the port: • **Authenticated**—The client has been authenticated through the RADIUS server or has been permitted access through server fail fallback. • **Authenticating**—The client is authenticating through the RADIUS server. • **Connecting**—Switch is attempting to contact the RADIUS server. • **Initialize**—The interface link is down. • **Held**—An action has been triggered through server fail fallback during a RADIUS server timeout. A supplicant is denied access, permitted access through a specified VLAN, or maintains the authenticated state granted to it before the RADIUS server timeout occurred.
MAC address	The MAC address of the connected client on the interface.
User	Users connected to the captive portal interface.

clear captive-portal interface

```
user@switch> clear captive-portal interface
ge-0/0/3.0
```

clear captive-portal interface

```
user@switch> clear captive-portal interface
Captive Portal Information:
Interface      State          MAC address          User
ge-0/0/3.0     Authenticated  00:03:47:e1:ba:b9    aclallow
ge-0/0/5.0     Connecting
ge-0/0/7.0     Connecting
ge-0/0/9.0     Connecting
```

clear captive-portal mac-address

```
user@switch> clear captive-portal mac-address 00:03:47:e1:ba:b9
```

This command has no output.

clear captive-portal firewall

```
user@switch> clear captive-portal firewall
```

This command has no output.

clear dot1x

Syntax **clear dot1x**
(interface (all | [*interface-names*]) | mac-address [*mac-addresses*])

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Reset the authentication state of a port. When you reset a port, reauthentication on the port is also triggered. The switch sends out a multicast message on the port to restart the authentication of all connected supplicants. If a MAC address is reset, then the switch sends out a unicast message to that specific MAC address to restart authentication.

If a supplicant is sending traffic when the **clear dot1x interface** command is issued, the authenticator immediately initiates reauthenticataion. This process happens very quickly, and it may seem that reauthentication did not occur. To verify that reauthentication has happened, issue the operational mode command **show dot1x interface detail**. The value for **Reauthentication due** and **Reauthentication interval** will be about the same.

Options all—(Optional) Clears all ports, or specific ports or specific MAC addresses.

interface *interface-names*—(Optional) Resets the authentication state of all supplicants connected to the specified ports (when the port is an authenticator) or for itself (when the port is a supplicant).

mac-address *mac-addresses*—Resets the authentication state only for the specified MAC addresses.

Required Privilege Level view

Related Documentation

- **show dot1x on page 2477**
- Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on a J-EX Series Switch on page 2290
- Filtering 802.1X Supplicants Using RADIUS Server Attributes on page 2340

List of Sample Output **clear dot1x interface on page 2468**
clear dot1x mac-address on page 2468

clear dot1x interface

```
user@switch> clear dot1x interface ge-1/0/0 ge-2/0/0 ge-2/0/0 ge5/0/0]
```

clear dot1x mac-address

```
user@switch> clear dot1x mac—address 00:04:ae:cd:23:5f
```

clear lldp neighbors

Syntax
clear lldp neighbors
<interface *interface*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Clear the learned remote neighbor information on all or selected interfaces.

Options
none—Clear the remote neighbor information on all interfaces.

interface *interface*—(Optional) Clear the remote neighbor information from one or more selected interfaces.

Required Privilege Level
view

Related Documentation
- **show lldp on page 2489**
- Configuring LLDP (CLI Procedure) on page 2344
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261

List of Sample Output
clear lldp neighbors on page 2469
clear lldp neighbors interface ge-0/1/1.0 on page 2469

clear lldp neighbors
```
user@switch> clear lldp neighbors
```

clear lldp neighbors interface ge-0/1/1.0
```
user@switch> clear lldp neighbors interface ge-0/1/1.0
```

clear lldp statistics

Syntax
clear lldp statistics
<interface *interface*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Clear LLDP statistics on one or more interfaces.

Options
none—Clears LLDP statistics on all interfaces.

interface *interface-names*—(Optional) Clear LLDP statistics on one or more interfaces.

Required Privilege Level
view

Related Documentation
- Configuring LLDP (CLI Procedure) on page 2344
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261

List of Sample Output
clear lldp statistics on page 2470
clear lldp statistics interface ge-0/1/1.0 on page 2470

clear lldp statistics
```
user@switch> clear lldp statistics
```

clear lldp statistics interface ge-0/1/1.0
```
user@switch> clear lldp statistics interface ge-0/1/1.0
```

show captive-portal authentication-failed-users

Syntax **show captive-portal authentication-failed-users**

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display the users that have failed captive portal authentication.

Required Privilege Level view

Related Documentation

- **show captive-portal interface on page 2474**
- **show captive-portal firewall on page 2472**
- **clear captive-portal on page 2466**
- Example: Setting Up Captive Portal Authentication on a J-EX Series Switch on page 2323
- Configuring Captive Portal Authentication (CLI Procedure) on page 2350

List of Sample Output **show captive-portal authentication-failed-users on page 2471**

Output Fields Table 305 on page 2471 lists the output fields for the **show captive-portal authentication-failed-users** command. Output fields are listed in the approximate order in which they appear.

Table 305: show captive-portal authentication-failed-users Output Fields

Field Name	Field Description	Level of Output
Interface	The MAC address configured to bypass captive portal authentication.	**all**
MAC address	The MAC address configured statically on the interface.	**all**
User	Name of the user that has failed captive portal authentication.	**all**

show captive-portal authentication-failed-users

```
user@switch> show captive-portal authentication-failed-users

Interface     MAC address          User
ge-0/0/10.0    00:00:00:10:00:02    md5user02
```

show captive-portal firewall

Syntax

show captive-portal firewall
<brief | detail>
<*interface-name*>
<*interface-name* detail>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display information about the firewall filters for each user that is authenticated on each captive portal interface.

Options

none—Display all the firewall filters on all captive portal interfaces.

brief | detail—(Optional) Display the specified level of output.

interface-name—(Optional) Display all the terms of the firewall filters for the specified interface.

interface-name detail—(Optional) Display all of the terms of the firewall filters for the specified interface.

Required Privilege Level

view

Related Documentation

- **show captive-portal authentication-failed-users on page 2471**
- **show captive-portal interface on page 2474**
- **clear captive-portal on page 2466**
- Example: Setting Up Captive Portal Authentication on a J-EX Series Switch on page 2323
- Configuring Captive Portal Authentication (CLI Procedure) on page 2350

List of Sample Output

show captive-portal firewall brief on page 2472
show captive-portal firewall ge-0/0/10.0 on page 2472
show captive-portal firewall on page 2473

Output Fields

Output fields for the **show captive-portal firewall** command include any action modifier specified in firewall filters except policers. Policers are not supported in the terms of the internally generated dynamic firewall filters that are created when multiple supplicants authenticate on 802.1X-enabled interfaces.

show captive-portal firewall brief

```
user@switch> show captive-portal firewall brief
Captive Portal Information:
Interface      State           MAC address          User
ge-0/0/1.0     Connecting
ge-0/0/10.0    Connecting      00:30:48:8c:66:bd    No User
```

show captive-portal firewall ge-0/0/10.0

```
user@switch> show captive-portal firewall ge-0/0/10.0
Filter name: dot1x_ge-0/0/10
Counters:
```

```
Name                          Bytes              Packets
dot1x_ge-0/0/10_CP_arp         7616                119
dot1x_ge-0/0/10_CP_dhcp           0                  0
dot1x_ge-0/0/10_CP_http           0                  0
dot1x_ge-0/0/10_CP_https          0                  0
dot1x_ge-0/0/10_CP_t_dns          0                  0
dot1x_ge-0/0/10_CP_u_dns          0                  0
```

show captive-portal firewall

```
user@switch> show captive-portal firewall
Filter name: dot1x_ge-0/0/0
Counters:
Name                                Bytes              Packets
dot1x_ge-0/0/0_CP_arp                  0                    0
dot1x_ge-0/0/0_CP_dhcp                 0                    0
dot1x_ge-0/0/0_CP_http                 0                    0
dot1x_ge-0/0/0_CP_https                0                    0
dot1x_ge-0/0/0_CP_t_dns                0                    0
dot1x_ge-0/0/0_CP_u_dns                0                    0
Filter name: dot1x_ge-0/0/1
Counters:
Name                                Bytes              Packets
dot1x_ge-0/0/1_CP_arp                  0                    0
dot1x_ge-0/0/1_CP_dhcp                 0                    0
dot1x_ge-0/0/1_CP_http                 0                    0
dot1x_ge-0/0/1_CP_https                0                    0
dot1x_ge-0/0/1_CP_t_dns                0                    0
dot1x_ge-0/0/1_CP_u_dns                0                    0
Filter name: dot1x_ge-0/0/10
Counters:
Name                                Bytes              Packets
dot1x_ge-0/0/10_CP_arp              7616                  119
dot1x_ge-0/0/10_CP_dhcp                0                    0
dot1x_ge-0/0/10_CP_http                0                    0
dot1x_ge-0/0/10_CP_https               0                    0
dot1x_ge-0/0/10_CP_t_dns               0                    0
dot1x_ge-0/0/10_CP_u_dns               0                    0
Filter name: dot1x_ge-0/0/11
```

show captive-portal interface

Syntax
show captive-portal interface
<interface-name>
detail

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display the current operational state of all captive portal interfaces with the list of connected users and the configured values of captive portal attributes on the interfaces.

Options
none—Display all captive portal interfaces.

interface-name—(Optional) Display the state for the specified captive portal interface and lists the MAC address and user names of any clients authenticated on the interface.

interface-name detail—(Optional) Displays the configured values of captive portal attributes on the specified captive portal interface.

Required Privilege Level
view

Related Documentation
- **show captive-portal authentication-failed-users on page 2471**
- **show captive-portal firewall on page 2472**
- **captive-portal on page 2385**
- **clear captive-portal on page 2466**
- Example: Setting Up Captive Portal Authentication on a J-EX Series Switch on page 2323
- Configuring Captive Portal Authentication (CLI Procedure) on page 2350

List of Sample Output
show captive-portal interface on page 2475
show captive-portal interface detail on page 2476

Output Fields
Table 306 on page 2474 lists the output fields for the **show captive-portal interface** command. Output fields are listed in the approximate order in which they appear.

Table 306: show captive-portal interface Output Fields

Field Name	Field Description	Level of Output
Interface	Interface on which captive portal has been configured.	All levels

Table 306: show captive-portal interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
State	The state of the interface: • **Authenticated**—The client has been authenticated through the RADIUS server or has been permitted access through server fail fallback. • **Authenticating**—The client is authenticating through the RADIUS server. • **Connecting**—Switch is attempting to contact the RADIUS server. • **Initialize**—The interface link is down. • **Held**—An action has been triggered through server fail fallback during a RADIUS server timeout. A supplicant is denied access, permitted access through a specified VLAN, or maintains the authenticated state granted to it before the RADIUS server timeout occurred.	All levels
MAC address	The MAC address of the connected client on the interface..	brief
User	Users connected to the captive portal interface.	brief
Supplicant mode	Mode used to authenticate clients—multiple, single, or single-supplicant.	detail
Number of retries	Number of times the user can attempt to submit authentication information.	detail
Quiet period	Time, in seconds, after a user exceeds the maximum number of retries before they can attempt to authenticate.	detail
Configured CP session timeout	Time, in seconds, that a client can be idle before the session expires.	detail
Server timeout	Time, in seconds, that an interface will wait for a reply when relaying a response from the client to the authentication server before timing out and invoking the server-fail action.	detail
Number of connected supplicants	Number of users connecting through the captive portal interface. Information for each user includes: • **Supplicant**—User name and MAC address. • **Operational state**—See State (above). • **Dynamic CP session timeout**—Timeout value dynamically downloaded from the RADIUS server for this user, if any. • **CP Session expiration due in**—Time remaining in session.	detail

show captive-portal interface

```
user@switch> show captive-portal interface
Captive Portal Information:
Interface      State           MAC address          User
```

```
ge-0/0/1.0     Connecting
ge-0/0/10.0    Connecting      00:30:48:8c:66:bd    No User
```

show captive-portal interface detail

```
user@switch> show captive-portal interface detail
ge-0/0/1.0
  Supplicant mode: Multiple
  Number of retries: 10
  Quiet period: 60 seconds
  Configured CP session timeout: 3600 seconds
  Server timeout: 15 seconds
  Number of connected supplicants: 0
ge-0/0/10.0
  Supplicant mode: Multiple
  Number of retries: 10
  Quiet period: 60 seconds
  Configured CP session timeout: 3600 seconds
  Server timeout: 15 seconds
  Number of connected supplicants: 1
      Supplicant: No User, 00:30:48:8c:66:bd
          Operational state: Connecting
          Dynamic CP Session Timeout: 0 seconds
          CP Session Expiration due in: 0 seconds
```

show dot1x

Syntax
show dot1x
<brief | detail>
<interface [*interface-names*]>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display the current operational state of all ports with the list of connected users.

Options none—Display information for all authenticator ports.

brief | detail—(Optional) Display the specified level of output.

interface *interface-names*—Display information for the specified port with a list of connected supplicants.

Required Privilege Level view

Related Documentation

- **clear dot1x on page 2468**
- Example: Setting Up 802.1X for Single Supplicant or Multiple Supplicant Configurations on a J-EX Series Switch on page 2290
- Example: Configuring 802.1X Authentication Options When the RADIUS Server is Unavailable to a J-EX Series Switch on page 2271
- Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch on page 2267
- Example: Configuring MAC RADIUS Authentication on a J-EX Series Switch on page 2286
- Example: Setting Up VoIP with 802.1X and LLDP-MED on a J-EX Series Switch on page 2302
- Configuring 802.1X RADIUS Accounting (CLI Procedure) on page 2339
- Filtering 802.1X Supplicants Using RADIUS Server Attributes on page 2340
- Verifying 802.1X Authentication on page 2356

List of Sample Output
show dot1x interface brief on page 2480
show dot1x interface detail on page 2480

Output Fields Table 307 on page 2477 lists the output fields for the **show dot1x** command. Output fields are listed in the approximate order in which they appear.

Table 307: show dot1x Output Fields

Field Name	Field Description	Level of Output
Interface	Name of a port.	All levels
MAC address	The MAC address of the connected supplicant on the port.	All levels

Table 307: show dot1x Output Fields *(continued)*

Field Name	Field Description	Level of Output
Role	The 802.1X authentication role of the interface. When 802.1X is enabled on an interface, the role is **Authenticator**. As **Authenticator**, the interface blocks LAN access until a supplicant is authenticated through 802.1X or MAC RADIUS authentication.	**brief**, **detail**
State	The state of the port: • **Authenticated**—The supplicant has been authenticated through the RADIUS server or has been permitted access through server fail fallback. • **Authenticating**—The supplicant is authenticating through the RADIUS server. • **Held**—An action has been triggered through server fail fallback during a RADIUS server timeout. A supplicant is denied access, permitted access through a specified VLAN, or maintains the authenticated state granted to it before the RADIUS server timeout occurred.	**brief**
Administrative state	The administrative state of the port: • **auto**—Traffic is allowed through the port based on the authentication result. (Default) • **force-authorize**—All traffic flows through the port irrespective of the authentication result. This state is not allowed on an interface whose VLAN membership has been set to **dynamic**. • **force-unauthorize**—All traffic drops on the port irrespective of the authentication result. This state is not allowed on an interface whose VLAN membership has been set to **dynamic**.	**detail**
Supplicant	The mode for the supplicant: • **single**—Authenticates only the first supplicant. All other supplicants who connect later to the port are allowed full access without any further authentication. They effectively “piggyback” on the first supplicant’s authentication. • **single-secure**—Allows only one supplicant to connect to the port. No other supplicant is allowed to connect until the first supplicant logs out. • **multiple**—Allows multiple supplicants to connect to the port. Each supplicant is authenticated individually.	**detail**
Quiet period	The number of seconds the port remains in the wait state following a failed authentication exchange with the supplicant before reattempting the authentication. The default value is 60 seconds. The range is 0 through 65,535 seconds.	**detail**
Transmit period	The number of seconds the port waits before retransmitting the initial EAPOL PDUs to the supplicant. The default value is 30 seconds. The range is 1 through 65,535 seconds.	**detail**
MAC radius	MAC RADIUS authentication: • **enabled**—The switch sends an EAPOL request to the connecting host to attempt 802.1X authentication and if the connecting host is unresponsive, the switch tries to authenticate using the MAC address. • **disabled**—The default. The switch will not attempt to authenticate the MAC address of the connecting host.	**detail**

Table 307: show dot1x Output Fields *(continued)*

Field Name	Field Description	Level of Output
MAC radius restrict	The authentication method is restricted to MAC RADIUS only. 802.1X authentication is not enabled.	**detail**
Reauthentication	The reauthentication state: • **disable**—Periodic reauthentication of the client is disabled. • **interval**—Sets the periodic reauthentication time interval. The default value is 3600 seconds. The range is 1 through 65,535 seconds.	**detail**
Supplicant timeout	The number of seconds the port waits for a response when relaying a request from the authentication server to the supplicant before resending the request. The default value is 30 seconds. The range is 1 through 60 seconds.	**detail**
Server timeout	The number of seconds the port waits for a reply when relaying a response from the supplicant to the authentication server before timing out. The default value is 30 seconds. The range is 1 through 60 seconds.	**detail**
Maximum EAPOL requests	The maximum number of retransmission times of an EAPOL request packet to the supplicant before the authentication session times out. The default value is 2. The range is 1 through 10.	**detail**
Number of clients bypassed because of authentication	The number of non-802.1X clients granted access to the LAN by means of static MAC bypass. The following fields are displayed: • **Client**—MAC address of the client. • **vlan** —The name of the VLAN to which the client is connected.	**detail**
Guest VLAN member	The VLAN to which a supplicant is connected when the supplicant is authenticated using a guest VLAN. If a guest VLAN is not configured on the interface, this field displays <**not configured**>.	**detail**
Number of connected supplicants	The number of supplicants connected to a port.	**detail**
Supplicant	The user name and MAC address of the connected supplicant.	**detail**

Table 307: show dot1x Output Fields *(continued)*

Field Name	Field Description	Level of Output
Authentication method	The 802.1X authentication method used for a supplicant: • **Guest VLAN**—A supplicant is connected to the LAN through the guest VLAN. • **MAC Radius**—A nonresponsive host is authenticated based on its MAC address. The MAC address is configured as permitted on the RADIUS server, the RADIUS server lets the switch know that the MAC address is a permitted address, and the switch opens LAN access to the nonresponsive host on the interface to which it is connected. • **Radius**—A supplicant is configured on the RADIUS server, the RADIUS server communicates this to the switch, and the switch opens LAN access on the interface to which the supplicant is connected. • **Server-fail deny**—If the RADIUS servers time out, all supplicants are denied access to the LAN, preventing traffic from flowing from the supplicant through the interface. This is the default. • **Server-fail permit**—When the RADIUS server is unavailable, a supplicant is still permitted access to the LAN as if the supplicant had been successfully authenticated by the RADIUS server. • **Server-fail use-cache**—If the RADIUS servers time out during reauthentication, previously authenticated supplicants are reauthenticated, but new supplicants are denied LAN access. • **Server-fail VLAN**—A supplicant is configured to be moved to a specified VLAN if the RADIUS server is unavailable to reauthenticate the supplicant. (The VLAN must already exist on the switch.)	**detail**
Authenticated VLAN	The VLAN to which the supplicant is connected.	**detail**
Dynamic filter	User policy filter sent by the RADIUS server.	**detail**
Session Reauth interval	The configured reauthentication interval.	**detail**
Reauthentication due in	The number of seconds in which reauthentication will occur again for the connected supplicant.	**detail**

show dot1x interface brief

```
user@switch> show dot1x interface [ge-0/0/1 ge-0/0/2 ge0/0/3] brief

Interface Role          State          MAC address
--------- ----          -----          ------------------
ge-0/0/1  Authenticator Authenticated  00:a0:d2:18:1a:c8
                        Authenticating 00:a0:e5:32:97:af
ge-0/0/2  Authenticator Connecting     -
ge-0/0/3  Supplicant    Authenticated  00:a6:55:f2:94:ae
```

show dot1x interface detail

```
user@switch> show dot1x interface ge-0/0/16.0 detail

 ge-0/0/16.0
  Role: Authenticator
  Administrative state: Auto
  Supplicant mode: Single
  Number of retries: 3
  Quiet period: 60 seconds
  Transmit period: 30 seconds
```

```
Mac Radius: Enabled
Mac Radius Strict: Disabled
Reauthentication: Enabled
Reauthentication interval: 40 seconds
Supplicant timeout: 30 seconds
Server timeout: 30 seconds
Maximum EAPOL requests: 1
Guest VLAN member: <not configured>
Number of connected supplicants: 1
  Supplicant: abc, 00:30:48:8C:66:BD
    Operational state: Authenticated
    Authentication method: Radius
    Authenticated VLAN: v200
    Reauthentication due in 17 seconds
```

show dot1x authentication-failed-users

Syntax show dot1x authentication-failed-users

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Displays supplicants (users) that have failed 802.1X authentication.

Required Privilege Level view

Related Documentation

- **clear dot1x on page 2468**
- Example: Configuring Static MAC Bypass of Authentication on a J-EX Series Switch on page 2281
- Configuring 802.1X Interface Settings (CLI Procedure) on page 2331

List of Sample Output **show dot1x authentication-failed-users on page 2482**

Output Fields Table 308 on page 2482 lists the output fields for the **show dot1x authentication-failed-users** command. Output fields are listed in the approximate order in which they appear.

Table 308: show dot1x authentication-failed-users Output Fields

Field Name	Field Description	Level of Output
Interface	The MAC address configured to bypass 802.1X authentication.	**all**
MAC address	The MAC address configured statically on the interface.	**all**
User	The user that is configured on the RADIUS server and that has failed 802.1X authentication.	**all**

show dot1x authentication-failed-users

```
user@switch> show dot1x authentication-failed-users

Interface      MAC address          User
ge-0/0/0.0     00:00:00:10:00:02    md5user02
```

show dot1x firewall

Syntax show dot1x firewall <interface *interface-name*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Displays information about the firewall filters for each user or nonresponsive host that is authenticated on each 802.1X-enabled interface that is configured for multiple supplicants. For example, if the firewall filter is configured with a term for counters, the command shows the count for each user.

Options interface *interface-names*—(Optional) Display information for the specified interface.

Required Privilege Level view

Related Documentation

- **clear dot1x on page 2468**
- Example: Applying Firewall Filters to Multiple Supplicants on 802.1X-Enabled Interfaces on page 2318

List of Sample Output

show dot1x firewall on page 2483
show dot1x firewall on page 2483

Output Fields Output fields include any action modifier that is specified in firewall filters.

show dot1x firewall (Showing counter action)

```
user@switch> show dot1x firewall
Filter: dot1x-filter-ge-0/0/3
Counters
 counter1_dot1x_ge-0/0/3_user1   342
 counter1_dot1x_ge-0/0/3_user2   857
```

show dot1x firewall (Showing policer action)

```
user@switch> show dot1x firewall
Filter: dot1x_ge-0/0/0
Counters
 p1-t1   494946
```

show dot1x static-mac-address

Syntax show dot1x static-mac-address <(interface [*interface-name*])>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Displays all the static MAC addresses that are configured to bypass 802.1X authentication on the switch.

Options **interface** [***interface-name***]—(Optional) Display static MAC addresses for a specific interface.

Required Privilege Level view

Related Documentation

- **clear dot1x on page 2468**
- Example: Configuring Static MAC Bypass of Authentication on a J-EX Series Switch on page 2281
- Configuring 802.1X Interface Settings (CLI Procedure) on page 2331
- Understanding Static MAC Bypass of Authentication on J-EX Series Switches

List of Sample Output **show dot1x static-mac-address on page 2484**
show dot1x static-mac-address interface ge-0/0/0.1 on page 2484

Output Fields Table 309 on page 2484 lists the output fields for the **show dot1x static-mac-address** command. Output fields are listed in the approximate order in which they appear.

Table 309: show dot1x static-mac-address Output Fields

Field Name	Field Description	Level of Output
MAC address	The MAC address of the device that is configured to bypass 802.1X authentication.	**all**
VLAN-Assignment	The name of the VLAN to which the device is assigned.	**all**
Interface	The name of the interface on which authentication is bypassed for a given MAC address.	**all**

show dot1x static-mac-address

```
user@switch> show dot1x static-mac-address

MAC address           VLAN-Assignment           Interface
00:00:00:11:22:33
00:00:00:00:12:12                               ge-0/0/3.0
00:00:00:02:34:56      facilities               ge-0/0/1.0
```

show dot1x static-mac-address interface ge-0/0/0.1

```
user@switch> show dot1x static-mac-address interface ge-0/0/0.1

MAC address           VLAN-Assignment           Interface
```

```
00:00:00:12:24:12     support                 ge-0/0/1.0
00:00:00:72:30:58     support                 ge-0/0/1.0
```

show ethernet-switching interfaces

Syntax show ethernet-switching interfaces
<brief | detail | summary>
<interface *interface-name*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display information about switched Ethernet interfaces.

Options none—(Optional) Display brief information for Ethernet switching interfaces.

brief | detail | summary—(Optional) Display the specified level of output.

interface *interface-name*—(Optional) Display Ethernet switching information for a specific interface.

Required Privilege Level view

Related Documentation

- **show ethernet-switching mac-learning-log on page 1241**
- **show ethernet-switching table on page 1249**
- Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure) on page 2516

List of Sample Output

show ethernet-switching interfaces on page 2487
show ethernet-switching interfaces ge-0/0/15 brief on page 2488
show ethernet-switching interfaces ge-0/0/2 detail (Blocked by RTG rtggroup) on page 2488
show ethernet-switching interfaces ge-0/0/15 detail (Blocked by STP) on page 2488
show ethernet-switching interfaces ge-0/0/17 detail (Disabled by bpdu-control) on page 2488
show ethernet-switching interfaces detail (C-VLAN to S-VLAN Mapping) on page 2488

Output Fields Table 310 on page 2486 lists the output fields for the **show ethernet-switching interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 310: show ethernet-switching interfaces Output Fields

Field Name	Field Description	Level of Output
Interface	Name of a switching interface.	All levels
State	Interface state. Values are **up** and **down**.	none, **brief**, **detail**, **summary**
VLAN members	Name of a VLAN.	none, **brief**, **detail**, **summary**
Tag	Number of the 802.1Q-tag.	All levels

Table 310: show ethernet-switching interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
Tagging	Specifies whether the interface forwards 802.1Q-tagged or untagged traffic.	All levels
Blocking	The forwarding state of the interface: • **unblocked**—Traffic is forwarded on the interface. • **blocked**—Traffic is not being forwarded on the interface. • **Disabled by bpdu control**—The interface is disabled due to receiving BPDUs on a protected interface. If the **disable-timeout** statement has been included in the BPDU configuration, the interface automatically returns to service after the timer expires. • **blocked by RTG**—The specified redundant trunk group is disabled. • **blocked by STP**—The interface is disabled due to a spanning tree protocol error. • **MAC limit exceeded**—The interface is temporarily disabled due to a MAC limiting error. The disabled interface is automatically restored to service when the disable timeout expires. • **MAC move limit exceeded**—The interface is temporarily disabled due to a MAC move limiting error. The disabled interface is automatically restored to service when the disable timeout expires. • **Storm control in effect**—The interface is temporarily disabled due to a storm control error. The disabled interface is automatically restored to service when the disable timeout expires.	none, **brief**, **detail**, **summary**
Index	The VLAN index internal to the Junos OS.	**detail**
mapping	The C-VLAN to S-VLAN mapping information: • **dot1q-tunneled**—The interface maps all traffic to the S-VLAN (all-in-one bundling). • **native**—The interface maps untagged and priority tagged packets to the S-VLAN. • **push**—The interface maps packets to a firewall filter to an S-VLAN. • **policy-mapped**—The interface maps packets to a specifically defined S-VLAN. • *integer*—The interface maps packets to the specified S-VLAN.	**detail**

show ethernet-switching interfaces

```
user@switch> show ethernet-switching interfaces

Interface    State  VLAN members        Tag   Tagging  Blocking

ae0.0        up     default                   untagged unblocked
ge-0/0/2.0   up     vlan300           300     untagged blocked by RTG (rtggroup)
ge-0/0/3.0   up     default                            blocked by STP
ge-0/0/4.0   down   default                            MAC limit exceeded
ge-0/0/5.0   down   default                            MAC move limit exceeded
ge-0/0/6.0   down   default                            Storm control in effect
ge-0/0/7.0   down   default                            unblocked
ge-0/0/13.0  up     default                   untagged unblocked
ge-0/0/14.0  up     vlan100             100   tagged   unblocked
                    vlan200             200   tagged   unblocked
ge-0/0/15.0  up     vlan100             100   tagged   blocked by STP
                    vlan200             200   tagged   blocked by STP
```

```
ge-0/0/16.0  down    default                  untagged unblocked
ge-0/0/17.0  down    vlan100            100   tagged   Disabled by bpdu-control

                     vlan200            200   tagged   Disabled by bpdu-control
```

show ethernet-switching interfaces ge-0/0/15 brief

```
user@switch> show ethernet-switching interfaces ge-0/0/15 brief
Interface   State  VLAN members   Tag   Tagging        Blocking

ge-0/0/15.0  up     vlan100        100   tagged         blocked by STP
                    vlan200        200   tagged         blocked by STP
```

show ethernet-switching interfaces ge-0/0/2 detail (Blocked by RTG rtggroup)

```
user@switch> show ethernet-switching interfaces ge-0/0/2 detail

Interface: ge-0/0/2.0, Index: 65, State: up, Port mode: Access
VLAN membership:
    vlan300, 802.1Q Tag: 300, untagged, msti-id: 0, blocked by RTG(rtggroup)
Number of MACs learned on IFL: 0
```

show ethernet-switching interfaces ge-0/0/15 detail (Blocked by STP)

```
user@switch> show ethernet-switching interfaces ge-0/0/15 detail

Interface: ge-0/0/15.0, Index: 70, State: up, Port mode: Trunk
VLAN membership:
    vlan100, 802.1Q Tag: 100, tagged, msti-id: 0, blocked by STP
    vlan200, 802.1Q Tag: 200, tagged, msti-id: 0, blocked by STP

Number of MACs learned on IFL: 0
```

show ethernet-switching interfaces ge-0/0/17 detail (Disabled by bpdu-control)

```
user@switch> show ethernet-switching interfaces ge-0/0/17 detail

Interface: ge-0/0/17.0, Index: 71, State: down, Port mode: Trunk
VLAN membership:
    vlan100, 802.1Q Tag: 100, tagged, msti-id: 1, Disabled by bpdu-control
    vlan200, 802.1Q Tag: 200, tagged, msti-id: 2, Disabled by bpdu-control
Number of MACs learned on IFL: 0
```

show ethernet-switching interfaces detail (C-VLAN to S-VLAN Mapping)

```
user@switch>show ethernet-switching interfaces ge-0/0/6.0 detail
Interface: ge-0/0/6.0, Index: 73, State: up, Port mode: Access
VLAN membership:
    map, 802.1Q Tag: 134, Mapped Tag: native, push, dot1q-tunneled, unblocked
    map, 802.1Q Tag: 134, Mapped Tag: 20, push, dot1q-tunneled, unblocked
```

show lldp

Syntax
show lldp
<detail>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display information about Link Layer Discovery Protocol (LLDP) and Link Level Discovery Protocol–Media Endpoint Discovery (LLDP-MED) configuration and capabilities on the switch. LLDP and LLDP-MED are used to learn about and to distribute device information on network links.

Options
none—Display LLDP information for all interfaces.

detail—(Optional) Display detailed LLDP information for all interfaces.

Required Privilege Level
view

Related Documentation
- Configuring LLDP (CLI Procedure) on page 2344
- Configuring LLDP-MED (CLI Procedure) on page 2346
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261

List of Sample Output
show lldp on page 2492
show lldp (detail) on page 2492

Output Fields
Table 311 on page 2489 lists the output fields for the **show lldp** command. Output fields are listed in the approximate order in which they appear.

Table 311: show lldp Output Fields

Field Name	Field Description	Level of Output
LLDP	LLDP operating state. The state can be **enabled** or **disabled**. NOTE: If a VLAN that has been configured for untagged packets on an interface also has Layer 2 protocol tunneling (L2PT) enabled for LLDP, the LLDP operating state for that interface is displayed as **disabled**.	All levels
Advertisement interval	Frequency, in seconds, at which LLDP advertisements are sent. This value is set by the **advertisement-interval** configuration statement.	All levels
Transmit delay	Delay between two successive LLDP advertisements. The delay is set to 2 seconds.	All levels
Hold timer	Multiplier used in combination with the **advertisement-interval** value to determine the length of time LLDP information is held before it is discarded. This value is set by the **hold-multiplier** configuration statement.	All levels

Table 311: show lldp Output Fields *(continued)*

Field Name	Field Description	Level of Output
Notification interval	How often LLDP trap notifications are generated as a result of LLDP database changes. If the interval value is 0, LLDP trap notifications of database changes are disabled. This value is set by the **lldp-configuration-notification-interval** configuration statement.	All levels
Config Trap Interval	How often LLDP trap notifications are generated as a result of changes in topology—for example, when an endpoint connects or disconnects. If the interval value is 0, LLDP trap notifications of topology changes are disabled. This value is set by the **ptopo-configuration-trap-interval** configuration statement.	All levels
Connection Hold timer	Amount of time the system maintains dynamic topology entries. This value is set by the **ptopo-configuration-maximum-hold-time** configuration statement.	All levels
LLDP-MED	LLDP-MED operating state. The state can be **enabled** or **disabled**.	All levels
LLDP-MED fast start count	Number of advertisements sent from a switch to a device, such as a VoIP telephone, when the device is first detected by the switch. These increased advertisements are temporary. After a device and a switch exchange information and can communicate, advertisements are reduced to one per second. This value is set by the **fast-start** configuration statement.	All levels
Interface	Name of the interface for which LLDP configuration information is being reported.	All levels
Parent Interface	Name of the aggregated Ethernet interface, if any, to which the interface belongs.	All levels
LLDP	LLDP operating state. The state can be **enabled** or **disabled**.	All levels
LLDP-MED	LLDP–MED operating state. The state can be **enabled** or **disabled**.	All levels
Neighbor count	Total number of new LLDP neighbors detected since the last switch reboot.	**detail**
Interface	Name of the interface that is advertising VLAN information.	All levels
Vlan-id	VLAN tag associated with the interface sending LLDP frames. If the interface is not a member of a VLAN, the VLAN ID is advertised as 0.	**detail**
Vlan-name	VLAN name associated with the VLAN ID.	**detail**

Table 311: show lldp Output Fields *(continued)*

Field Name	Field Description	Level of Output
LLDP basic TLVs supported	Basic TLVs supported on the switch: • **Chassis identifier**—TLV that advertises the MAC address associated with the local system. • **Port identifier**—TLV that advertises the port identification for the specified port in the local system. • **Port description**—TLV that advertises the user-configured port description. • **System name**—TLV that advertises the user-configured name of the local system. • **System description**—TLV that advertises the system description containing information about the software and current image running on the system. This information is taken from the software and is not configurable. • **System capabilities**—TLV that advertises the primary functions performed by the system—for example, bridge or router. • **Management address**—TLV that advertises the IP management address of the local system.	**detail**
Supported LLDP 802 TLVs	802.3 TLVs supported on the switch: • **Power via MDI**—TLV that advertises MDI power support, PSE power pair, and power class information. • **Link aggregation**—TLV that advertises if the interface is aggregated and its aggregated interface ID. • **Maximum frame size**—TLV that advertises the maximum transmission unit (MTU) of the interface sending LLDP frames. • **Port VLAN tag**—TLV that advertises the VLAN tag configured on the interface. • **Port VLAN name**—TLV that advertises the VLAN name configured on the interface.	**detail**

Supported LLDP MED TLVs	LLDP-MED TLVs supported on the switch:
	• **LLDP MED capabilities**—TLV that advertises the primary function of the port. The capabilities values range from 0 through 15:
	  ◦ **0**—Capabilities
	  ◦ **1**—Network Policy
	  ◦ **2**—Location Identification
	  ◦ **3**—Extended Power via MDI-PSE
	  ◦ **4**—Inventory
	  ◦ **5–15**—Reserved
	• **Network policy**—TLV that advertises the port VLAN configuration and associated Layer 2 and Layer 3 attributes. Attributes include the policy identifier, application types—such as voice or streaming video—802.1Q VLAN tagging, and 802.1p priority bits and DiffServ code points.
	• **Endpoint location**—TLV that advertises the physical location of the endpoint.
	• **Extended power Via MDI**—TLV that advertises the power type, power source, power priority, and power value of the port. It is the responsibility of the PSE device (network connectivity device) to advertise the power priority on a port.
		detail

show lldp

```
user@switch> show lldp

LLDP                     : Enabled
Advertisement interval   : 30 seconds
Transmit delay           : 2 seconds
Hold timer               : 4 seconds
Notification interval    : 0 Second(s)
Config Trap Interval     : 0 seconds
Connection Hold timer    : 300 seconds

LLDP MED                 : Disabled
MED fast start count     : 3 Packets


Interface      Parent Interface    LLDP        LLDP-MED
all            -                   Enabled     -
me0.0          -                   Disabled    -
```

show lldp (detail)

```
user@switch> show lldp detail

LLDP                     : Enabled
Advertisement interval   : 30 seconds
Transmit delay           : 2 seconds
Hold timer               : 4 seconds
Notification interval    : 0 Second(s)
Config Trap Interval     : 0 seconds
Connection Hold timer    : 300 seconds

LLDP MED                 : Disabled
MED fast start count     : 3 Packets


Interface      Parent Interface    LLDP        LLDP-MED    Neighbor count
all            -                   Enabled     -           8
me0.0          -                   Disabled    -           0

Interface      Parent Interface    Vlan-id     Vlan-name
xe-3/0/0.0     ae31.0              100         v100
xe-3/0/0.0     ae31.0              101         v101
xe-3/0/0.0     ae31.0              4000        v4000
xe-3/0/1.0     ae31.0              100         v100
xe-3/0/1.0     ae31.0              101         v101
xe-3/0/1.0     ae31.0              4000        v4000
xe-3/0/2.0     ae31.0              100         v100
xe-3/0/2.0     ae31.0              101         v101
xe-3/0/2.0     ae31.0              4000        v4000

LLDP basic TLVs supported:
Chassis identifier, Port identifier, Port description, System name, System
description, System capabilities, Management address.

Supported LLDP 802 TLVs:
Power via MDI, Link aggregation, Maximum frame size, Port VLAN tag, Port
VLAN name.

Supported LLDP MED TLVs:
LLDP MED capabilities, Network policy, Endpoint location, Extended power
Via MDI.
```

show lldp local-information

Syntax show lldp local-information

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Displays the information that the switch provides in Link Layer Discovery Protocol (LLDP) advertisements to its neighbors.

Required Privilege Level view

Related Documentation

- Configuring LLDP (CLI Procedure) on page 2344
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261

List of Sample Output **show lldp local-information on page 2494**

Output Fields Table 312 on page 2493 lists the output fields for the **show lldp local-information** command. Output fields are listed in the approximate order in which they appear.

Table 312: show lldp local-information Output Fields

Field Name	Field Description
LLDP Local Information details	Information about the local system (the switch): • **Chassis ID**—MAC address associated with the switch. • **System name**—User-configured name of the switch. • **System descr**—System description containing information about the switch model and the current software image running on the switch. This information is taken from the software and is not configurable.
System Capabilities	Capabilities (such as **bridge** or **router**) that are supported or enabled on the system.
Management Information	Details of the management information: **Port Name**, **Port Address** (such as **10.204.34.35**), **Address Type** (such as **ipv4** or **ipv6**), **Port ID** (SNMP interface index), **Subtype**, and **Interface Subtype**.
Interface Name	Name of the local interface.
Parent Interface	Name of the aggregated Ethernet interface, if any, to which the local interface belongs.
Interface ID	SNMP interface index.
Interface description	User-configured port description.
Status	Administrative status of the interface: either **up** or **down**.
Tunneling	Status of tunneling on the interface: either **enabled** or **disabled**.

show lldp local-information

```
user@switch> show lldp local-information

LLDP Local Information details

Chassis ID   : 00:1d:b5:aa:b9:f0
System name  : switch
System descr : Juniper Networks, Inc. ex8208 , version 10.3I0 [builder] Build
               date: 2010-03-24 12:38:30 UTC

System Capabilities
    Supported       : Bridge Router
    Enabled         : Bridge Router

Management Information
    Port Name       : -
    Port Address    : 10.93.54.6
    Address Type    : IPv4
    Port ID         : 34
    Port ID Subtype : local(7)
    Port Subtype    : ifIndex(1)

Interface name Parent Interface Interface ID Interface description Status Tunneling
me0.0          -                34           -                     Down   Disabled
xe-3/0/0.0     ae31.0           769          xe-3/0/0.0            Up     Disabled
xe-3/0/1.0     ae31.0           770          xe-3/0/1.0            Up     Disabled
xe-3/0/2.0     ae31.0           771          xe-3/0/2.0            Up     Disabled
xe-3/0/3.0     ae31.0           772          xe-3/0/3.0            Up     Disabled
xe-3/0/4.0     ae31.0           577          xe-3/0/4.0            Up     Disabled
xe-3/0/5.0     ae31.0           578          xe-3/0/5.0            Up     Disabled
xe-3/0/6.0     ae31.0           579          xe-3/0/6.0            Up     Disabled
xe-3/0/7.0     ae31.0           581          xe-3/0/7.0            Up     Disabled
```

show lldp neighbors

Syntax

show lldp neighbors
<interface *interface*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display the information about neighboring devices learned by the switch by using the Link Layer Discovery Protocol (LLDP).

Options

none—Display LLDP neighbor information for all interfaces.

interface *interface*—(Optional) Display LLDP neighbor information for a selected interface.

Required Privilege Level

view

Related Documentation

- Configuring LLDP (CLI Procedure) on page 2344
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261

List of Sample Output

show lldp neighbors on page 2497
show lldp neighbors interface xe-3/0/4.0 on page 2497
show lldp neighbors interface (for a VoIP Avaya Telephone with LLDP-MED Support) on page 2498

Output Fields

Table 313 on page 2495 lists the output fields for the **show lldp neighbors** command. Output fields are listed in the approximate order in which they appear.

Table 313: show lldp neighbors Output Fields

Field Name	Field Description
Local Interface	List of local interfaces for which neighbor information is available.
Parent Interface	List of aggregated Ethernet interfaces, if any, to which the local interfaces belong.
Chassis ID	List of chassis identifiers for neighbors.
Port info	List of port information gathered from neighbors. This could be the port identifier or port description.
System name	List of system names gathered from neighbors.
LLDP Neighbor Information	Information about both the local system (the switch) and a neighbor system on the interface (appears when the **interface** option is used).
Local Information	Information about the local system (appears when the **interface** option is used).
Index	Local interface index (appears when the **interface** option is used).

Table 313: show lldp neighbors Output Fields *(continued)*

Field Name	Field Description
Time mark	Date and timestamp of information (appears when the **interface** option is used).
Time to live	Number of seconds for which this information is valid (appears when the **interface** option is used).
Local Interface	Name of the local physical interface (appears when the **interface** option is used).
Parent Interface	Name of the aggregated Ethernet interface, if any, to which the interface belongs (appears when the **interface** option is used).
Local Port ID	Local interface SNMP index (appears when the interface option is used).
Ageout Count	Number of times the complete set of information advertised by the neighbor has been deleted from LLDP neighbor information maintained by the local system because the information timeliness interval has expired (appears when the interface option is used).
Neighbor Information	Information about a neighbor system on the interface (appears when the **interface** option is used).
Chassis type	Type of chassis identifier supplied, such as **MAC address** (appears when the **interface** option is used).
Chassis ID	Chassis identifier of the chassis type listed (appears when the **interface** option is used).
Port type	Type of port identifier supplied, such as **locally assigned** (appears when the **interface** option is used).
Port ID	Port identifier of the port type listed (appears when the **interface** option is used).
Port description	Port description (appears when the **interface** option is used).
System name	Name supplied by the system on the interface (appears when the **interface** option is used).
System Description	Description supplied by the system on the interface (appears when the **interface** option is used).
System capabilities	Capabilities (such as **Bridge**, **Router**, and **Telephone**) that are supported or enabled by the system on the interface (appears when the **interface** option is used).
Management Info	Details of management information: **Address Type** (such as **ipv4** or **ipv6**), **Address** (such as **10.204.34.35**), **Port ID**, **Subtype**, **Interface Subtype**, and organization identifier (**OID**) (appears when the **interface** option is used).

Table 313: show lldp neighbors Output Fields *(continued)*

Field Name	Field Description
Media Info	Additional details about the endpoint device appear when a device that supports LLDP-MED is attached to the interface. The specific details depend upon the capabilities of the device. Details may include: **Media endpoint class** (such as Class 3 for communication devices such as IP phones), **MED Hardware revision**, **MED Firmware revision**, **MED Software revision**, **MED Serial number**, **MED Manufacturer name**, **MED Model name**.
Organization Info	One or more entries listing remote information by organizationally unique identifier (**OUI**), **Subtype**, **Index**, and **Info** (appears when the **interface** option is used).

show lldp neighbors

```
user@switch> show lldp neighbors

Local Interface   Parent Interface   Chassis Id          Port info    System Name
xe-3/0/4.0        ae31.0             b0:c6:9a:63:80:40   xe-0/0/0.0   newyork31
xe-3/0/5.0        ae31.0             b0:c6:9a:63:80:40   xe-0/0/1.0   newyork31
xe-3/0/6.0        ae31.0             b0:c6:9a:63:80:40   xe-0/0/2.0   newyork31
xe-3/0/7.0        ae31.0             b0:c6:9a:63:80:40   xe-0/0/3.0   newyork31
xe-3/0/0.0        ae31.0             b0:c6:9a:63:80:40   xe-0/1/0.0   newyork31
xe-3/0/1.0        ae31.0             b0:c6:9a:63:80:40   xe-0/1/1.0   newyork31
xe-3/0/2.0        ae31.0             b0:c6:9a:63:80:40   xe-0/1/2.0   newyork31
xe-3/0/3.0        ae31.0             b0:c6:9a:63:80:40   xe-0/1/3.0   newyork31
```

show lldp neighbors interface xe-3/0/4.0

```
user@switch>show lldp neighbors interface xe-3/0/4.0

LLDP Neighbor Information:
Local Information:
Index: 488 Time to live: 120 Time mark: Tue Mar 30 23:33:28 2010 Age: 30 secs
Local Interface    : xe-3/0/4.0
Parent Interface   : ae31.0
Local Port ID      : 577
Ageout Count       : 10

Neighbour Information:
Chassis type       : Mac address
Chassis ID         : b0:c6:9a:63:80:40
Port type          : Locally assigned
Port ID            : 503
Port description   : xe-0/0/0.0
System name        : newyork31

System Description : Juniper Networks, Inc. ex4500-40f , version 10.2I0 Build
date: 2010-03-26 00:17:34 UTC

System capabilities
        Supported  : Bridge Router
        Enabled    : Bridge Router

Management Info
        Type              : IPv4
        Address           : 10.10.200.84
        Port ID           : 34
        Subtype           : 1
        Interface Subtype : 2
```

```
        OID              : 1.3.6.1.2.1.31.1.1.1.1.34

Organization Info
       OUI      : 0.18.15
       Subtype  : 1
       Index    : 1
       Info     : 0000010000

Organization Info
       OUI      : 0.18.15
       Subtype  : 3
       Index    : 2
       Info     : 0300000207

Organization Info
       OUI      : 0.18.15
       Subtype  : 4
       Index    : 3
       Info     : 05EA

Organization Info
       OUI      : 0.18.15
       Subtype  : 1
       Index    : 4
       Info     : 444530323039343530333438

Organization Info
       OUI      : 0.18.15
       Subtype  : 3
       Index    : 5
       Info     : 00640476313030

Organization Info
       OUI      : 0.18.15
       Subtype  : 3
       Index    : 6
       Info     : 00650476313031

Organization Info
       OUI      : 0.18.15
       Subtype  : 3
       Index    : 7
       Info     : 0FA0057634303030
```

show lldp neighbors interface (for a VoIP Avaya Telephone with LLDP-MED Support)

```
user@switch>show lldp neighbors interface ge-0/0/0.0

LLDP Neighbor Information:
Local Information:
Index: 20 Time to live: 120 Time mark: Thu Apr 15 22:26:22 2010 Age: 16 secs
Local Interface    : ge-0/0/0.0
Parent Interface   : -
Local Port ID      : 517
Ageout Count       : 0

Neighbour Information:
Chassis type       : Network address
Chassis ID         : 0.0.0.0
Port type          : Mac address
Port ID            : 00:04:0d:fc:55:48
System name        : AVAFC5548
```

```
System capabilities
        Supported  : Bridge Telephone
        Enabled    : Bridge

Management Info
        Type              : IPv4
        Address           : 0.0.0.0
        Port ID           : 1
        Subtype           : 1
        Interface Subtype : 2
        OID               : 1.3.6.1.2.1.31.1.1.1.1.1
Media endpoint class: Class III Device

MED Hardware revision : 4610D01A
MED Firmware revision : b10d01b2_9.bin
MED Software revision : a10d01b2_9.bin
MED Serial number     : 07N510103424
MED Manufacturer name : Avaya
MED Model name        : 4610

Organization Info
       OUI      : 0.18.15
       Subtype  : 1
       Index    : 1
       Info     : 036CA00010

Organization Info
       OUI      : 0.18.15
       Subtype  : 1
       Index    : 2
       Info     : 002303

Organization Info
       OUI      : 0.18.15
       Subtype  : 2
       Index    : 3
       Info     : 014001AE

Organization Info
       OUI      : 0.18.15
       Subtype  : 5
       Index    : 4
       Info     : 3436313044303141

Organization Info
       OUI      : 0.18.15
       Subtype  : 6
       Index    : 5
       Info     : 62313064303162325F392E62696E

Organization Info
       OUI      : 0.18.15
       Subtype  : 7
       Index    : 6
       Info     : 613130643031623 25F392E62696E

Organization Info
       OUI      : 0.18.15
       Subtype  : 8
       Index    : 7
       Info     : 30374E353130313033343234
```

```
Organization Info
       OUI      : 0.18.15
       Subtype  : 9
       Index    : 8
       Info     : 4176617961

Organization Info
       OUI      : 0.18.15
       Subtype  : 10
       Index    : 9
       Info     : 34363130

Organization Info
       OUI      : 0.18.15
       Subtype  : 1
       Index    : 10
       Info     : 000028003C

Organization Info
       OUI      : 0.18.15
       Subtype  : 3
       Index    : 11
       Info     : 00000000

Organization Info
       OUI      : 0.18.15
       Subtype  : 4
       Index    : 12
       Info     : 000000000000000000000000

Organization Info
       OUI      : 0.18.15
       Subtype  : 5
       Index    : 13
       Info     : 00000000

Organization Info
       OUI      : 0.18.15
       Subtype  : 6
       Index    : 14
       Info     : 00000000

Organization Info
       OUI      : 0.18.15
       Subtype  : 7
       Index    : 15
       Info     : 01
```

show lldp remote-global-statistics

Syntax show lldp remote-global-statistics

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display remote Link Layer Discovery Protocol (LLDP) global statistics.

Options This command has no options.

Required Privilege Level view

Related Documentation

- Configuring LLDP (CLI Procedure) on page 2344
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261

List of Sample Output **show lldp remote-global-statistics on page 2502**

Output Fields Table 314 on page 2501 describes the output fields for the **show lldp remote-global-statistics** command. Output fields are listed in the approximate order in which they appear.

Table 314: show lldp remote-global-statistics Output Fields

Field Name	Field Description
LLDP Remote Database Table Counters	Information about remote database table counters.
LastchangeTime	Time elapsed between LLDP agent startup and the last change to the remote database table information.
Inserts	Number of insertions made in the remote database table.
Deletes	Number of deletions made in the remote database table.
Drops	Number of LLDP frames dropped from the remote database table because of errors.
Ageouts	Number of remote database table entries that have aged out of the table.

show lldp remote-global-statistics

```
user@host> show lldp remote-global-statistics
user@host> show lldp remote-global-statistics
   LLDP Remote Database Table Counters
   LastchangeTime          Inserts    Deletes    Drops    Ageouts
   00:00:76 (76 sec)       192        0          0        0
```

show lldp statistics

Syntax **show lldp statistics**
<interface *interface*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display LLDP statistics for all interfaces or for the specified interface.

Options none—Display LLDP statistics for all interfaces.

interface *interface*—(Optional) Display LLDP statistics for the specified interface.

Required Privilege Level view

Related Documentation
- Configuring LLDP (CLI Procedure) on page 2344
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261

List of Sample Output
show lldp statistics on page 2504
show lldp statistics interface xe-3/0/0.0 on page 2504

Output Fields Table 315 on page 2503 lists the output fields for the **show lldp statistics** command. Output fields are listed in the approximate order in which they appear.

Table 315: show lldp statistics Output Fields

Field Name	Field Description
Interface	Name of the interface.
Parent Interface	Name of the aggregated Ethernet interface, if any, to which the interface belongs. NOTE: Because LLDP packets are transmitted and received on member interfaces only, statistics are available only for the member interfaces, not for the aggregated interface.
Received	Total number of LLDP frames received on an interface.
Unknown TLVs	Number of unrecognized LLDP TLVs received on an interface.
With Errors	Number of invalid LLDP TLVs received on an interface.
Discarded	Number of LLDP TLVs received and then discarded on an interface.
Transmitted	Total number of LLDP frames that were transmitted on an interface.
Untransmitted	Total number of LLDP frames that were untransmitted on an interface.

show lldp statistics

```
user@switch> show lldp statistics

Interface    Parent Interface  Received  Unknown TLVs  With Errors
xe-3/0/0.0   ae31.0            1564      0             0
xe-3/0/1.0   ae31.0            1564      0             0
xe-3/0/2.0   ae31.0            1565      0             0
xe-3/0/3.0   ae31.0            1566      0             0
xe-3/0/4.0   ae31.0            1598      0             0
xe-3/0/5.0   ae31.0            1598      0             0
xe-3/0/6.0   ae31.0            1596      0             0
xe-3/0/7.0   ae31.0            1597      0             0
xe-5/0/6.0   -                 0         0             0
xe-5/0/7.0   -                 0         0             0


Discarded TLVs  Transmitted  Untransmitted
0               3044         1
0               3044         1
0               3044         1
0               3044         1
0               3075         1
0               3075         1
0               3075         1
0               3075         1
0               17312        0
0               17312        0
```

show lldp statistics interface xe-3/0/0.0

```
user@switch> show lldp statistics interface xe-3/0/0.0

Interface    Parent Interface  Received  Unknown TLVs  With Errors
xe-3/0/0.0   ae31.0            1566      0             0

Discarded TLVs  Transmitted  Untransmitted
0               3046         1
```

show network-access aaa statistics accounting

Syntax **show network-access aaa statistics accounting**

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display authentication, authorization, and accounting (AAA) accounting statistics.

Required Privilege Level view

Related Documentation

- **accounting-server on page 2373**
- **accounting-stop-on-access-deny on page 2374**
- Configuring 802.1X RADIUS Accounting (CLI Procedure) on page 2339

List of Sample Output **show network-access aaa statistics accounting on page 2505**

Output Fields Table 316 on page 2505 lists the output fields for the **show network-access aaa statistics accounting** command. Output fields are listed in the approximate order in which they appear.

Table 316: show network-access aaa statistics accounting Output Fields

Field Name	Field Description
Requests received	The number of accounting-request packets sent from a switch to a RADIUS accounting server.
Accounting Response failures	The number of accounting-response failure packets sent from the RADIUS accounting server to the switch.
Accounting Response Success	The number of accounting-response success packets sent from the RADIUS accounting server to the switch.
Requests timedout	The number of requests-timedout packets sent from the RADIUS accounting server to the switch.

show network-access aaa statistics accounting

```
user@switch> show network-access aaa statistics accounting
Accounting module statistics
  Requests received: 1
  Accounting Response failures: 0
  Accounting Response Success: 1
  Requests timedout: 0
```

show network-access aaa statistics authentication

Syntax **show network-access aaa statistics authentication**

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display authentication, authorization, and accounting (AAA) authentication statistics.

Required Privilege Level view

Related Documentation
- **authentication-server on page 2383**
- Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch on page 2267

List of Sample Output **show network-access aaa statistics authentication on page 2506**

Output Fields Table 317 on page 2506 lists the output fields for the **show network-access aaa statistics authentication** command. Output fields are listed in the approximate order in which they appear.

Table 317: show network-access aaa statistics authentication Output Fields

Field Name	Field Description
Requests received	The number of authentication requests received by the switch.
Accepts	The number of authentication accepts received by the RADIUS server.
Rejects	The number authentication rejects sent by the RADIUS server.
Challenges	The number of authentication challenges sent by the RADIUS server.

show network-access aaa statistics authentication

```
user@switch> show network-access aaa statistics authentication
Authentication module statistics
  Requests received: 2
  Accepts: 1
  Rejects: 0
  Challenges: 1
```

show network-access aaa statistics dynamic-requests

Syntax **show network-access aaa statistics dynamic-requests;**

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display authentication, authorization, and accounting (AAA) authentication statistics for disconnects.

Required Privilege Level view

Related Documentation
- **authentication-server on page 2383**
- Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch on page 2267

List of Sample Output **show network-access aaa statistics authentication on page 2507**

Output Fields Table 318 on page 2507 lists the output fields for the **show network-access aaa statistics dynamic-requests** command. Output fields are listed in the approximate order in which they appear.

Table 318: show network-access aaa statistics dynamic-requests Output Fields

Field Name	Field Description
Requests received	The number of dynamic requests received by the RADIUS server.
Processed successfully	The number of dynamic requests successfully processed by the RADIUS server.
Errors during processing	The number of errors that occurred while the RADIUS server was processing the dynamic request.
Silently dropped	The number of silently dropped requests.

show network-access aaa statistics authentication

```
user@switch> show network-access aaa statistics dynamic-requests
Dynamic-requests module statistics
  Requests received: 0
  Processed successfully: 0
  Errors during processing: 0
  Silently dropped: 0
```

PART 18

Rate Limiting

CHAPTER 87

Rate Limiting Overview

Understanding Storm Control on J-EX Series Switches

A traffic storm is generated when messages are broadcast on a network and each message prompts a receiving node to respond by broadcasting its own messages on the network. This, in turn, prompts further responses, creating a snowball effect. The LAN is suddenly flooded with packets, creating unnecessary traffic that leads to poor network performance or even a complete loss of network service. Storm control enables the switch to monitor traffic levels and to drop broadcast and unknown unicast packets when a specified traffic level—called the *storm control level*—is exceeded, thus preventing packets from proliferating and degrading the LAN. As an alternative to having the switch drop packets, you can configure it to shut down interfaces or temporarily disable interfaces (see the **action-shutdown** statement or the **port-error-disable** statement) when the storm control level is exceeded.

The factory default configuration enables storm control on all switch interfaces, with the storm control level set to 80 percent of the combined broadcast and unknown unicast streams. You can change the storm control level for an interface by specifying a bandwidth value for the combined broadcast and unknown unicast traffic streams. You can also selectively disable storm control on the broadcast stream or on the unknown unicast stream.

Broadcast, multicast, and unicast packets are part of normal LAN operation, so to recognize a storm, you must be able to identify when traffic has reached a level that is abnormal for your LAN. Suspect a storm when operations begin timing out and network response times slow down. As more packets flood the LAN, network users might be unable to access servers or e-mail.

Monitor the level of broadcast and unknown unicast traffic in the LAN when it is operating normally. Use this data as a benchmark to determine when traffic levels are too high. Then configure storm control to set the level at which you want to drop broadcast traffic, unknown unicast traffic, or both.

NOTE: When you configure storm control bandwidth on an aggregated Ethernet interface, the storm control level for each member of the aggregated Ethernet interface is set to that bandwidth. For example, if you configure a storm control bandwidth of 15000 Kbps on ae1, and ae1 has two members, ge-0/0/0 and ge-0/0/1, each member has a storm control level of 15000 Kbps. Thus, the storm control level on ae1 allows a traffic rate of up to 30000 Kbps of combined broadcast and unknown unicast traffic.

Related Documentation

- Example: Configuring Storm Control to Prevent Network Outages on J-EX Series Switches on page 2513
- Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure) on page 2516

Understanding Unknown Unicast Forwarding on J-EX Series Switches

Unknown unicast traffic consists of unicast packets with unknown destination MAC addresses. By default, the switch floods these unicast packets that are traveling in a VLAN to all interfaces that are members of the VLAN. Forwarding this type of traffic to interfaces on the switch can trigger a security issue. The LAN is suddenly flooded with packets, creating unnecessary traffic that leads to poor network performance or even a complete loss of network service. This is known as a traffic storm.

To prevent a storm, you can disable the flooding of unknown unicast packets to all interfaces by configuring one VLAN or all VLANs to forward and unknown unicast traffic to a specific trunk interface. This channels the unknown unicast traffic to a single interface.

Related Documentation

- Understanding Storm Control on J-EX Series Switches on page 2511
- Example: Configuring Storm Control to Prevent Network Outages on J-EX Series Switches on page 2513
- Configuring Unknown Unicast Forwarding (CLI Procedure) on page 2515

CHAPTER 88

Example: Rate Limiting Configuration

Example: Configuring Storm Control to Prevent Network Outages on J-EX Series Switches

Storm control enables you to prevent network outages caused by broadcast storms on the LAN. You can configure storm control on the J-EX Series switch to rate limit broadcast traffic and unknown unicast traffic at a specified level and to drop packets when the specified traffic level is exceeded, thus preventing packets from proliferating and degrading the LAN.

This example shows how to configure storm control on a single J-EX Series switch:

Requirements

This example uses the following hardware and software components:

- One J-EX Series switch

Overview and Topology

A storm is generated when messages are broadcast on a network and each message prompts a receiving node to respond by broadcasting its own messages on the network. This, in turn, prompts further responses, creating a snowball effect and resulting in a broadcast storm that can cause network outages.

You can use storm control to prevent broadcast storms by specifying the amount, also known as the storm control level, of broadcast traffic and unknown unicast traffic to be allowed on an interface. You specify the storm control level as the traffic rate in kilobits per second of the combined broadcast and unknown unicast streams.

NOTE: The factory default configuration enables storm control on all interfaces at 80 percent of the combined broadcast and unknown unicast streams.

Storm control monitors the incoming broadcast traffic and unknown unicast traffic and compares it with the level that you specify. If broadcast traffic and unknown unicast traffic exceed the specified level, the switch drops packets for the controlled traffic types. As an alternative to having the switch drop packets, you can configure it to shut down interfaces or temporarily disable interfaces (see the**action-shutdown** statement or the **port-error-disable** statement) when the storm control level is exceeded.

The topology used in this example consists of one J-EX Series switch with 24 ports. The switch is connected to various network devices. This example shows how to configure the storm control level on interface **ge-0/0/0** by setting the level to a traffic rate of 15000 Kbps, based on the traffic rate of the combined broadcast and unknown unicast streams. If broadcast traffic and unknown unicast traffic exceeds this level, the switch drops packets for the controlled traffic types to prevent a network outage.

Configuration

CLI Quick Configuration

To quickly configure storm control based on the traffic rate in kilobits per second of the combined broadcast and unknown unicast streams, copy the following command and paste it into the switch terminal window:

```
[edit]
set ethernet-switching-options storm-control interface ge-0/0/0 bandwidth 15000
```

Step-by-Step Procedure

To configure storm control:

1. Specify the traffic rate in kilobits per second of the combined broadcast and unknown unicast streams on a specific interface:

```
[edit ethernet-switching-options]
user@switch# set storm-control interface ge-0/0/0 bandwidth 15000
```

Results

Display the results of the configuration:

```
[edit ethernet-switching-options]
user@switch# show storm-control
interface ge-0/0/0.0 {
   bandwidth 15000;
}
```

Related Documentation

- Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure) on page 2516
- Understanding Storm Control on J-EX Series Switches on page 2511

CHAPTER 89

Configuring Rate Limiting

Configuring Unknown Unicast Forwarding (CLI Procedure)

Unknown unicast traffic consists of packets with unknown destination MAC addresses. By default, the switch floods these packets to all interfaces associated with a VLAN. Forwarding such traffic to interfaces on the switch can create a security issue.

To prevent flooding unknown unicast traffic across the switch, configure unknown unicast forwarding to direct all unknown unicast packets within a VLAN out to a specific trunk interface. From there, the destination MAC address can be learned and added to the Ethernet switching table. You can configure each VLAN to divert unknown unicast traffic to different trunk interfaces or use one trunk interface for multiple VLANs.

To configure unknown unicast forwarding options using the CLI:

NOTE: Before you can configure unknown unicast forwarding within a VLAN, you must first configure that VLAN.

1. Configure unknown unicast forwarding for a specific VLAN (here, the VLAN name is **employee**):

```
[edit ethernet-switching-options]
user@switch# set unknown-unicast-forwarding vlan employee
```

2. Specify the trunk interface to which all unknown unicast traffic will be forwarded:

```
[edit ethernet-switching-options ]
user@switch# set unknown-unicast-forwarding vlan employee interface ge-0/0/3.0
```

Related Documentation

- Example: Configuring Storm Control to Prevent Network Outages on J-EX Series Switches on page 2513
- Configuring VLANs for J-EX Series Switches (CLI Procedure) on page 1136
- Configuring VLANs for J-EX Series Switches (J-Web Procedure) on page 1133

- Verifying That Unknown Unicast Packets Are Forwarded to a Trunk Interface on page 2517
- Understanding Unknown Unicast Forwarding on J-EX Series Switches on page 2512
- Understanding Storm Control on J-EX Series Switches on page 2511

Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure)

An Ethernet switching access interface on a J-EX Series switch might shut down or be disabled as a result of one of the following port-security or storm-control configurations:

- MAC limiting—**mac-limit** statement is configured with action **shutdown**.
- MAC move limiting—**mac-move-limit** statement is configured with action **shutdown**.
- Storm control—**storm-control** statement is configured with the action **shutdown**.

You can configure the switch to automatically restore the disabled interfaces to service after a specified period of time. Autorecovery applies to all the interfaces that have been disabled due to MAC limiting, MAC move limiting, or storm control errors.

NOTE: You must specify the disable timeout value for the interfaces to recover automatically. There is no default disable timeout. If you do not specify a timeout value, you need to use the clear ethernet-switching port-error command to clear the errors and restore the interfaces or the specified interface to service.

To configure autorecovery from the disabled state due to MAC limiting, MAC move limiting, or storm control shutdown actions:

```
[edit ethernet-switching-options]
user@switch# set port-error-disable disable-timeout 60
```

Related Documentation

- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Configuring MAC Limiting (CLI Procedure) on page 2635
- Example: Configuring Storm Control to Prevent Network Outages on J-EX Series Switches on page 2513
- Understanding MAC Limiting and MAC Move Limiting for Port Security on J-EX Series Switches on page 2557
- Understanding Storm Control on J-EX Series Switches on page 2511

CHAPTER 90

Verifying Rate Limiting Configuration

Verifying That Unknown Unicast Packets Are Forwarded to a Trunk Interface

Purpose Verify that a VLAN is forwarding all unknown unicast packets (those with unknown destination MAC addresses) to a single trunk interface instead of flooding unknown unicast packets across all interfaces that are members of the same VLAN.

Action Display the forwarding interface for unknown unicast packets for a VLAN (here, the VLAN name is **v1**):

```
user@switch> show configuration ethernet-switching-options

unknown-unicast-forwarding {
    vlan v1 {
        interface ge-0/0/7.0;
    }
}
```

Display the Ethernet switching table:

```
user@switch> show ethernet-switching table vlan v1
Ethernet-switching table: 3 unicast entries
  VLAN              MAC address       Type         Age Interfaces
  v1                *                 Flood          - All-members
  v1                00:01:09:00:00:00 Learn         24 ge-0/0/7.0
  v1                00:11:09:00:01:00 Learn         37 ge-0/0/3.0
```

Meaning The sample output from the **show configuration ethernet-switching-options** command shows that the unknown unicast forwarding interface for VLAN **v1** is interface **ge-0/0/7**. The **show ethernet-switching table** command shows that an unknown unicast packet is received on interface **ge-0/0/3** with the destination MAC address (DMAC) **00:01:09:00:00:00** and the source MAC address (SMAC) of **00:11:09:00:01:00**. This shows that the SMAC of the packet is learned in the normal way (through the interface **ge-0/0/3.0**), while the DMAC is learned on interface **ge-0/0/7**.

Related Documentation

- Configuring Unknown Unicast Forwarding (CLI Procedure) on page 2515

Verifying That the Port Error Disable Setting Is Working Correctly

Purpose Verify that the port error disable setting is working as expected on MAC limited, MAC move limited and rate-limited interfaces on a J-EX Series switch.

Action Display information about interfaces:

```
user@switch> show ethernet-switching interfaces
Interface   State    VLAN members           Blocking
ge-0/0/0.0  up       T1122                  unblocked
ge-0/0/1.0  down     default                MAC limit exceeded
ge-0/0/2.0  down     default                MAC move limit exceeded
ge-0/0/3.0  down     default                Storm control in effect
ge-0/0/4.0  down     default                unblocked
ge-0/0/5.0  down     default                unblocked
ge-0/0/6.0  down     default                unblocked
ge-0/0/7.0  down     default                unblocked
ge-0/0/8.0  down     default                unblocked
ge-0/0/9.0  up       T111                   unblocked
ge-0/0/10.0 down     default                unblocked
ge-0/0/11.0 down     default                unblocked
ge-0/0/12.0 down     default                unblocked
ge-0/0/13.0 down     default                unblocked
ge-0/0/14.0 down     default                unblocked
ge-0/0/15.0 down     default                unblocked
ge-0/0/16.0 down     default                unblocked
ge-0/0/17.0 down     default                unblocked
ge-0/0/18.0 down     default                unblocked
ge-0/0/19.0 up       T111                   unblocked
ge-0/1/0.0  down     default                unblocked
ge-0/1/1.0  down     default                unblocked
ge-0/1/2.0  down     default                unblocked
ge-0/1/3.0  down     default                unblocked
```

Meaning The sample output from the **show ethernet-switching interfaces** command shows that three of the down interfaces specify the reason that the interface is disabled:

- **MAC limit exceeded**—The interface is temporarily disabled due to a **mac-limit** error. The disabled interface is automatically restored to service when the **disable-timeout** expires.
- **MAC move limit exceeded**—The interface is temporarily disabled due to a **mac-move-limit** error. The disabled interface is automatically restored to service when the **disable-timeout** expires.
- **Storm control in efffect** —The interface is temporarily disabled due to a **storm-control** error. The disabled interface is automatically restored to service when the **disable-timeout** expires.

Related Documentation

- Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure) on page 2516

CHAPTER 91

Configuration Statements for Rate Limiting

[edit ethernet-switching-options] Configuration Statement Hierarchy

```
ethernet-switching-options {
  analyzer {
    name {
      loss-priority priority;
      ratio number;
      input {
        ingress {
          interface (all | interface-name);
          vlan (vlan-id | vlan-name);
        }
        egress {
          interface (all | interface-name);
        }
      output {
        interface interface-name;
        vlan (vlan-id | vlan-name);
      }
    }
  }
  bpdu-block {
    disable-timeout timeout;
    interface (all | [interface-name]);
  }
  dot1q-tunneling {
    ether-type (0x8100 | 0x88a8 | 0x9100);
  }
  interfaces interface-name {
    no-mac-learning;
  }
  mac-notification {
    notification-interval seconds;
  }
  mac-table-aging-time seconds;
  port-error-disable {
    disable-timeout timeout;
```

```
}
redundant-trunk-group {
  group-name name {
    interface interface-name <primary>;
  }
}
secure-access-port {
  dhcp-snooping-file {
    location local_pathname | remote_URL;
    timeout seconds;
    write-interval seconds;
  }
  interface (all | interface-name) {
    allowed-mac {
      mac-address-list;
    }
    (dhcp-trusted | no-dhcp-trusted );
    mac-limit limit action action;
    no-allowed-mac-log;
    static-ip  ip-address {
      vlan  vlan-name;
      mac mac-address;
    }
  }
  vlan (all | vlan-name) {
    (arp-inspection | no-arp-inspection );
    dhcp-option82 {
      circuit-id {
        prefix hostname;
        use-interface-description;
        use-vlan-id;
      }
      remote-id {
        prefix hostname | mac | none;
        use-interface-description;
        use-string string;
      }
      vendor-id [string];
    }
    (examine-dhcp | no-examine-dhcp );
    (ip-source-guard | no-ip-source-guard);
    mac-move-limit limit action action;
  }
}
storm-control {
  action-shutdown;
  interface (all | interface-name) {
    bandwidth bandwidth;
    no-broadcast;
    no-unknown-unicast;
  }
}
traceoptions {
  file filename <files number> <no-stamp> <replace> <size size> <world-readable |
    no-world-readable>;
  flag flag <disable>;
```

```
    }
    unknown-unicast-forwarding {
      vlan (all | vlan-name) {
        interface interface-name;
      }
    }
    voip {
      interface (all | [interface-name | access-ports]) {
        vlan vlan-name ;
        forwarding-class <assured-forwarding | best-effort | expedited-forwarding |
          network-control>;
      }
    }
  }
```

Related Documentation

- Understanding Port Mirroring on J-EX Series Switches on page 3245
- Port Security for J-EX Series Switches Overview on page 2545
- Understanding BPDU Protection for STP, RSTP, and MSTP on J-EX Series Switches on page 1278
- Understanding Redundant Trunk Links on J-EX Series Switches on page 1049
- Understanding Storm Control on J-EX Series Switches on page 2511
- Understanding 802.1X and VoIP on J-EX Series Switches on page 2263
- Understanding Q-in-Q Tunneling on J-EX Series Switches on page 1051
- Understanding Unknown Unicast Forwarding on J-EX Series Switches on page 2512
- Understanding MAC Notification on J-EX Series Switches on page 1060

action-shutdown

Syntax action-shutdown;

Hierarchy Level [edit **ethernet-switching-options storm-control**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Shut down or disable interfaces when the storm control level is exceeded, as follows:

- If you set both the **action-shutdown** and the **port-error-disable** statements, the interfaces are disabled temporarily and recover automatically when the disable timeout expires.
- If you set the **action-shutdown** statement and do not the specify the **port-error-disable** statement, the interfaces that are enabled for storm control are shut down when the storm control level is exceeded and they do not recover automatically from that port-error condition. You must issue the **clear ethernet-switching port-error** command to clear the port error and restore the interfaces to service.

Default The **action-shutdown** option is not enabled. When the storm control level is exceeded, the switch drops unknown unicast and broadcast messages on the specified interfaces.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **port-error-disable on page 2531**
- **disable-timeout on page 2524**
- clear ethernet-switching port-error
- Example: Configuring Storm Control to Prevent Network Outages on J-EX Series Switches on page 2513
- Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure) on page 2516
- Understanding Storm Control on J-EX Series Switches on page 2511

bandwidth

Syntax bandwidth *bandwidth*;

Hierarchy Level [edit **ethernet-switching-options storm-control interface** (all | *interface-name*)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the storm control level as the bandwidth in kilobits per second of the combined broadcast and unknown unicast streams.

NOTE: When you configure storm control bandwidth on an aggregated Ethernet interface, the storm control level for each member of the aggregated Ethernet interface is set to that bandwidth. For example, if you configure a storm control bandwidth of 15000 Kbps on ae1, and ae1 has two members, ge-0/0/0 and ge-0/0/1, each member has a storm control level of 15000 Kbps. Thus, the storm control level on ae1 allows a traffic rate of up to 30000 Kbps of combined broadcast and unknown unicast traffic.

Default If you omit the **bandwidth** statement when you configure storm control on an interface, the storm control level defaults to 80 percent of the combined broadcast and unknown unicast streams.

Options **bandwidth**—Traffic rate in kilobits per second of the combined broadcast and unknown unicast streams.

Range: 100 through 10000000

Default: None

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Storm Control to Prevent Network Outages on J-EX Series Switches on page 2513
- Understanding Storm Control on J-EX Series Switches on page 2511

disable-timeout

Syntax disable-timeout *timeout*;

Hierarchy Level [edit **ethernet-switching-options port-error-disable**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify how long the Ethernet-switching interfaces remain in a disabled state due to the MAC limiting, MAC move liming, or storm control errors.

Default The disable timeout is not enabled.

Options ***timeout*** —Amount of time, in seconds, that the disabled state remains in effect. The disabled interface is automatically restored to service when the specified timeout is reached.

Range: 10 through 3600 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Port Security (CLI Procedure) on page 2626
- Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure) on page 2516
- Example: Configuring Storm Control to Prevent Network Outages on J-EX Series Switches on page 2513

ethernet-switching-options

Syntax

```
ethernet-switching-options {
  analyzer {
    name {
      loss-priority priority;
      ratio number;
      input {
        ingress {
          interface (all | interface-name);
          vlan (vlan-id | vlan-name);
        }
        egress {
          interface (all | interface-name);
        }
      output {
        interface interface-name;
        vlan (vlan-id | vlan-name);
      }
    }
  }
  bpdu-block {
    disable-timeout timeout;
    interface (all | [interface-name]);
  }
  dot1q-tunneling {
    ether-type (0x8100 | 0x88a8 | 0x9100);
  }
  interfaces interface-name {
    no-mac-learning;
  }
  mac-notification {
    notification-interval seconds;
  }
  mac-table-aging-time seconds;
  port-error-disable {
    disable-timeout timeout;
  }
  redundant-trunk-group {
    group-name name {
      interface interface-name <primary>;
      interface interface-name;
    }
  }
  secure-access-port {
    dhcp-snooping-file {
      location local_pathname | remote_URL;
      timeout seconds;
      write-interval seconds;
    }
    interface (all | interface-name) {
      allowed-mac {
        mac-address-list;
      }
```

```
        (dhcp-trusted | no-dhcp-trusted);
        mac-limit limit action action;
        no-allowed-mac-log;
        static-ip ip-address {
          vlan  vlan-name;
          mac mac-address;
        }
      }
      vlan (all | vlan-name) {
        (arp-inspection | no-arp-inspection);
        dhcp-option82 {
          circuit-id {
            prefix hostname;
            use-interface-description;
            use-vlan-id;
          }
          remote-id {
            prefix hostname | mac | none;
            use-interface-description;
            use-string string;
          }
          vendor-id [string];
        }
        (examine-dhcp | no-examine-dhcp);
        (ip-source-guard | no-ip-source-guard);
        mac-move-limit limit action action;
      }
    }
    storm-control {
      action-shutdown;
      interface (all | interface-name) {
        bandwidth bandwidth;
        no-broadcast;
        no-unknown-unicast;
      }
    }
    traceoptions {
      file filename <files number> <no-stamp> <replace> <size size> <world-readable |
        no-world-readable>;
      flag flag <disable>;
    }
    unknown-unicast-forwarding {
      vlan (all | vlan-name) {
        interface interface-name;
      }
    }
    voip {
      interface (all | [interface-name | access-ports]) {
        vlan vlan-name ;
        forwarding-class <assured-forwarding | best-effort | expedited-forwarding |
          network-control>;
      }
    }
  }
```

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure Ethernet switching options.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation

- Understanding Port Mirroring on J-EX Series Switches on page 3245
- Port Security for J-EX Series Switches Overview on page 2545
- Understanding BPDU Protection for STP, RSTP, and MSTP on J-EX Series Switches on page 1278
- Understanding Redundant Trunk Links on J-EX Series Switches on page 1049
- Understanding Storm Control on J-EX Series Switches on page 2511
- Understanding 802.1X and VoIP on J-EX Series Switches on page 2263
- Understanding Q-in-Q Tunneling on J-EX Series Switches on page 1051
- Understanding Unknown Unicast Forwarding on J-EX Series Switches on page 2512
- Understanding MAC Notification on J-EX Series Switches on page 1060

interface

Syntax

```
interface (all | interface-name) {
    bandwidth bandwidth;
    no-broadcast;
    no-unknown-unicast;
}
```

Hierarchy Level [edit ethernet-switching-options storm-control]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable and configure storm control on all interfaces or on the specified interface.

If you do not include the **bandwidth** statement, the storm control level defaults to 80 percent of the combined broadcast and unknown unicast streams.

Default The factory default configuration enables storm control on all switch interfaces at the default level of 80 percent of the combined broadcast and unknown unicast streams.

Options **all**—All interfaces. The storm control settings configured with the **all** option affect only those interfaces that have not been individually configured for storm control.

interface-name—Name of an interface. The storm control settings configured with the ***interface-name*** option override any settings configured with the **all** option.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Storm Control to Prevent Network Outages on J-EX Series Switches on page 2513
- Understanding Storm Control on J-EX Series Switches on page 2511

interface

Syntax interface *interface-name*;

Hierarchy Level [edit **ethernet-switching-options unknown-unicast-forwarding vlan**(all|*vlan-name*)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the interface to which unknown unicast packets will be forwarded.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show vlans on page 1263**
- **show ethernet-switching table on page 1249**
- Configuring Unknown Unicast Forwarding (CLI Procedure) on page 2515
- Understanding Unknown Unicast Forwarding on J-EX Series Switches on page 2512

no-broadcast

Syntax no-broadcast;

Hierarchy Level [edit **ethernet-switching-options storm-control interface** (all | *interface-name*)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable storm control for broadcast traffic for the specified interface or for all interaces.

Default Storm control is enabled for both unknown unicast traffic and broadcast traffic.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Storm Control to Prevent Network Outages on J-EX Series Switches on page 2513
- Understanding Storm Control on J-EX Series Switches on page 2511

no-unknown-unicast

Syntax no-unknown-unicast;

Hierarchy Level [edit **ethernet-switching-options storm-control interface** (all | *interface-name*)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable storm control for unknown unicast traffic for the specified interface or for all interfaces.

Default Storm control is enabled for both unknown unicast traffic and broadcast traffic.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Storm Control to Prevent Network Outages on J-EX Series Switches on page 2513
- Understanding Storm Control on J-EX Series Switches on page 2511

port-error-disable

Syntax

```
port-error-disable {
  disable-timeout timeout ;
}
```

Hierarchy Level [edit **ethernet-switching-options**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable rather than block an interface when enforcing MAC limiting, MAC move limiting, and rate-limiting configuration options for shutting down the interface; and allow the interface to recovery automatically from the error condition after a specified period of time:

NOTE: The **port-error-disable** configuration does not apply to pre-existing error conditions. It impacts only error conditions that are detected after **port-error-disable** has been enabled and committed. To clear a pre-existing error condition and restore the interface to service, use the clear ethernet-switching port-error command.

- If you have enabled **mac-limit** with the **shutdown** option and enable **port-error-disable**, the switch disables (rather than shuts down) the interface when the MAC address limit is reached.
- If you have enabled **mac-move-limit** with the **shutdown** option and you enable **port-error-disable**, the switch disables (rather than shuts down) the interface when the maximum number of moves to a new interface is reached.
- If you have enabled **storm-control** with the **action-shutdown** option and you enable **port-error-disable**, the switch disables (rather than shuts down) the interface when broadcast traffic and unknown unicast traffic exceeds the specified levels.

Default Not enabled.

Required Privilege Level routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation

- Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure) on page 2516
- Configuring Port Security (CLI Procedure) on page 2626
- Example: Configuring Storm Control to Prevent Network Outages on J-EX Series Switches on page 2513

storm-control

Syntax

```
storm-control {
  action-shutdown;
  interface (all | interface-name) {
    bandwidth bandwidth;
    no-broadcast;
    no-unknown-unicast;
  }
}
```

Hierarchy Level [edit ethernet-switching-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure storm control on the switch.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Storm Control to Prevent Network Outages on J-EX Series Switches on page 2513
- Understanding Storm Control on J-EX Series Switches on page 2511

unknown-unicast-forwarding

Syntax

```
unknown-unicast-forwarding {
  vlan (all | vlan-name){
    interface interface-name;
  }
}
```

Hierarchy Level [edit ethernet-switching-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the switch to forward all unknown unicast packets in a VLAN or on all VLANs to a particular interface.

NOTE: Before you can configure unknown unicast forwarding within a VLAN, you must first configure that VLAN.

The remaining statements are explained separately.

Default Unknown unicast packets are flooded to all interfaces that belong to the same VLAN.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show vlans on page 1263**
- **show ethernet-switching table on page 1249**
- Configuring Unknown Unicast Forwarding (CLI Procedure) on page 2515
- Understanding Unknown Unicast Forwarding on J-EX Series Switches on page 2512

vlan

Syntax

```
vlan (all | vlan-name) {
  interface interface-name;
}
```

Hierarchy Level [edit **ethernet-switching-options unknown-unicast-forwarding**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify a VLAN from which unknown unicast packets will be forwarded or specify that the packets will be forwarded from all VLANS. Unknown unicast packets are forwarded from a VLAN to a specific trunk interface.

The **interface** statement is explained separately.

TIP: To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type ? after vlan or vlans in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.

Options all—All VLANs.

vlan-name—Name of a VLAN.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **show vlans on page 1263**
- **show ethernet-switching table on page 1249**
- Configuring Unknown Unicast Forwarding (CLI Procedure) on page 2515
- Verifying That Unknown Unicast Packets Are Forwarded to a Trunk Interface on page 2517
- Understanding Unknown Unicast Forwarding on J-EX Series Switches on page 2512

CHAPTER 92

Operational Mode Commands for Rate Limiting

show ethernet-switching interfaces

Syntax show ethernet-switching interfaces
<brief | detail | summary>
<interface *interface-name*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display information about switched Ethernet interfaces.

Options none—(Optional) Display brief information for Ethernet switching interfaces.

brief | detail | summary—(Optional) Display the specified level of output.

interface *interface-name*—(Optional) Display Ethernet switching information for a specific interface.

Required Privilege Level view

Related Documentation

- **show ethernet-switching mac-learning-log on page 1241**
- **show ethernet-switching table on page 1249**
- Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure) on page 2516

List of Sample Output

show ethernet-switching interfaces on page 2537
show ethernet-switching interfaces ge-0/0/15 brief on page 2538
show ethernet-switching interfaces ge-0/0/2 detail (Blocked by RTG rtggroup) on page 2538
show ethernet-switching interfaces ge-0/0/15 detail (Blocked by STP) on page 2538
show ethernet-switching interfaces ge-0/0/17 detail (Disabled by bpdu-control) on page 2538
show ethernet-switching interfaces detail (C-VLAN to S-VLAN Mapping) on page 2538

Output Fields Table 319 on page 2536 lists the output fields for the **show ethernet-switching interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 319: show ethernet-switching interfaces Output Fields

Field Name	Field Description	Level of Output
Interface	Name of a switching interface.	All levels
State	Interface state. Values are **up** and **down**.	none, **brief**, **detail**, **summary**
VLAN members	Name of a VLAN.	none, **brief**, **detail**, **summary**
Tag	Number of the 802.1Q-tag.	All levels

Table 319: show ethernet-switching interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
Tagging	Specifies whether the interface forwards 802.1Q-tagged or untagged traffic.	All levels
Blocking	The forwarding state of the interface: • **unblocked**—Traffic is forwarded on the interface. • **blocked**—Traffic is not being forwarded on the interface. • **Disabled by bpdu control**—The interface is disabled due to receiving BPDUs on a protected interface. If the **disable-timeout** statement has been included in the BPDU configuration, the interface automatically returns to service after the timer expires. • **blocked by RTG**—The specified redundant trunk group is disabled. • **blocked by STP**—The interface is disabled due to a spanning tree protocol error. • **MAC limit exceeded**—The interface is temporarily disabled due to a MAC limiting error. The disabled interface is automatically restored to service when the disable timeout expires. • **MAC move limit exceeded**—The interface is temporarily disabled due to a MAC move limiting error. The disabled interface is automatically restored to service when the disable timeout expires. • **Storm control in effect**—The interface is temporarily disabled due to a storm control error. The disabled interface is automatically restored to service when the disable timeout expires.	none, **brief**, **detail**, **summary**
Index	The VLAN index internal to the Junos OS.	**detail**
mapping	The C-VLAN to S-VLAN mapping information: • **dot1q-tunneled**—The interface maps all traffic to the S-VLAN (all-in-one bundling). • **native**—The interface maps untagged and priority tagged packets to the S-VLAN. • **push**—The interface maps packets to a firewall filter to an S-VLAN. • **policy-mapped**—The interface maps packets to a specifically defined S-VLAN. • *integer*—The interface maps packets to the specified S-VLAN.	**detail**

show ethernet-switching interfaces

```
user@switch> show ethernet-switching interfaces

Interface    State  VLAN members        Tag   Tagging  Blocking

ae0.0        up     default                   untagged unblocked
ge-0/0/2.0   up     vlan300           300     untagged blocked by RTG (rtggroup)
ge-0/0/3.0   up     default                            blocked by STP
ge-0/0/4.0   down   default                            MAC limit exceeded
ge-0/0/5.0   down   default                            MAC move limit exceeded
ge-0/0/6.0   down   default                            Storm control in effect
ge-0/0/7.0   down   default                            unblocked
ge-0/0/13.0  up     default                   untagged unblocked
ge-0/0/14.0  up     vlan100             100   tagged   unblocked
                    vlan200             200   tagged   unblocked
ge-0/0/15.0  up     vlan100             100   tagged   blocked by STP
                    vlan200             200   tagged   blocked by STP
```

```
ge-0/0/16.0  down    default                 untagged unblocked
ge-0/0/17.0  down    vlan100            100  tagged   Disabled by bpdu-control

                     vlan200            200  tagged   Disabled by bpdu-control
```

show ethernet-switching interfaces ge-0/0/15 brief

```
user@switch> show ethernet-switching interfaces ge-0/0/15 brief
Interface   State  VLAN members   Tag   Tagging        Blocking

ge-0/0/15.0  up     vlan100        100   tagged         blocked by STP
                    vlan200        200   tagged         blocked by STP
```

show ethernet-switching interfaces ge-0/0/2 detail (Blocked by RTG rtggroup)

```
user@switch> show ethernet-switching interfaces ge-0/0/2 detail

Interface: ge-0/0/2.0, Index: 65, State: up, Port mode: Access
VLAN membership:
    vlan300, 802.1Q Tag: 300, untagged, msti-id: 0, blocked by RTG(rtggroup)
Number of MACs learned on IFL: 0
```

show ethernet-switching interfaces ge-0/0/15 detail (Blocked by STP)

```
user@switch> show ethernet-switching interfaces ge-0/0/15 detail

Interface: ge-0/0/15.0, Index: 70, State: up, Port mode: Trunk
VLAN membership:
    vlan100, 802.1Q Tag: 100, tagged, msti-id: 0, blocked by STP
    vlan200, 802.1Q Tag: 200, tagged, msti-id: 0, blocked by STP

Number of MACs learned on IFL: 0
```

show ethernet-switching interfaces ge-0/0/17 detail (Disabled by bpdu-control)

```
user@switch> show ethernet-switching interfaces ge-0/0/17 detail

Interface: ge-0/0/17.0, Index: 71, State: down, Port mode: Trunk
VLAN membership:
    vlan100, 802.1Q Tag: 100, tagged, msti-id: 1, Disabled by bpdu-control
    vlan200, 802.1Q Tag: 200, tagged, msti-id: 2, Disabled by bpdu-control
Number of MACs learned on IFL: 0
```

show ethernet-switching interfaces detail (C-VLAN to S-VLAN Mapping)

```
user@switch>show ethernet-switching interfaces ge-0/0/6.0 detail
Interface: ge-0/0/6.0, Index: 73, State: up, Port mode: Access
VLAN membership:
    map, 802.1Q Tag: 134, Mapped Tag: native, push, dot1q-tunneled, unblocked
    map, 802.1Q Tag: 134, Mapped Tag: 20, push, dot1q-tunneled, unblocked
```

show ethernet-switching table

Syntax
show ethernet-switching table
<brief | detail | extensive | summary>
<interface *interface-name*>
<management-vlan>
<sort-by (*name* | *tag*)
<vlan (*vlan-name*)>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Displays the Ethernet switching table.

Options
none—(Optional) Display brief information about the Ethernet switching table.

brief | detail | extensive | summary—(Optional) Display the specified level of output.

management-vlan—(Optional) Display the Ethernet switching table for a management VLAN.

interface-name—(Optional) Display the Ethernet switching table for a specific interface.

sort-by (*name* | *tag*)—(Optional) Display VLANs in ascending order of VLAN IDs or VLAN names.

vlan *vlan-name*—(Optional) Display the Ethernet switching table for a specific VLAN.

Required Privilege Level
view

Related Documentation
- Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch on page 1063
- Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches on page 1070
- Example: Configure Automatic VLAN Administration Using GVRP on page 1087
- Example: Setting Up Q-in-Q Tunneling on J-EX Series Switches on page 1105

List of Sample Output
show ethernet-switching table on page 2540
show ethernet-switching table brief on page 2541
show ethernet-switching table detail on page 2541
show ethernet-switching table extensive on page 2542
show ethernet-switching table interface ge-0/0/1 on page 2542

Output Fields
Table 320 on page 2539 lists the output fields for the **show ethernet-switching table** command. Output fields are listed in the approximate order in which they appear.

Table 320: show ethernet-switching table Output Fields

Field Name	Field Description	Level of Output
VLAN	The name of a VLAN.	All levels

Table 320: show ethernet-switching table Output Fields *(continued)*

Field Name	Field Description	Level of Output
Tag	The VLAN ID tag name or number.	**extensive**
MAC or **MAC address**	The MAC address associated with the VLAN.	All levels
Type	The type of MAC address. Values are: • **static**—The MAC address is manually created. • **learn**—The MAC address is learned dynamically from a packet's source MAC address. • **flood**—The MAC address is unknown and flooded to all members.	All levels
Age	The time remaining before the entry ages out and is removed from the Ethernet switching table.	All levels
Interfaces	Interface associated with learned MAC addresses or **All-members** (flood entry).	All levels
Learned	For learned entries, the time which the entry was added to the Ethernet-switching table.	**detail**, **extensive**
Nexthop index	The nexthop index number.	**detail**, **extensive**

show ethernet-switching table

```
user@switch> show ethernet-switching table
Ethernet-switching table: 57 entries, 17 learned
  VLAN              MAC address       Type         Age Interfaces
  F2                *                 Flood          - All-members
  F2                00:00:05:00:00:03 Learn          0 ge-0/0/44.0
  F2                00:19:e2:50:7d:e0 Static         - Router
  Linux             *                 Flood          - All-members
  Linux             00:19:e2:50:7d:e0 Static         - Router
  Linux             00:30:48:90:54:89 Learn          0 ge-0/0/47.0
  T1                *                 Flood          - All-members
  T1                00:00:05:00:00:01 Learn          0 ge-0/0/46.0
  T1                00:00:5e:00:01:00 Static         - Router
  T1                00:19:e2:50:63:e0 Learn          0 ge-0/0/46.0
  T1                00:19:e2:50:7d:e0 Static         - Router
  T10               *                 Flood          - All-members
  T10               00:00:5e:00:01:09 Static         - Router
  T10               00:19:e2:50:63:e0 Learn          0 ge-0/0/46.0
  T10               00:19:e2:50:7d:e0 Static         - Router
  T111              *                 Flood          - All-members
  T111              00:19:e2:50:63:e0 Learn          0 ge-0/0/15.0
  T111              00:19:e2:50:7d:e0 Static         - Router
  T111              00:19:e2:50:ac:00 Learn          0 ge-0/0/15.0
  T2                *                 Flood          - All-members
  T2                00:00:5e:00:01:01 Static         - Router
  T2                00:19:e2:50:63:e0 Learn          0 ge-0/0/46.0
  T2                00:19:e2:50:7d:e0 Static         - Router
  T3                *                 Flood          - All-members
  T3                00:00:5e:00:01:02 Static         - Router
  T3                00:19:e2:50:63:e0 Learn          0 ge-0/0/46.0
  T3                00:19:e2:50:7d:e0 Static         - Router
  T4                *                 Flood          - All-members
```

```
  T4                00:00:5e:00:01:03 Static         - Router
  T4                00:19:e2:50:63:e0 Learn          0 ge-0/0/46.0
[output truncated]
```

show ethernet-switching table brief

```
user@switch> show ethernet-switching table brief
Ethernet-switching table: 57 entries, 17 learned
  VLAN              MAC address       Type          Age Interfaces
  F2                *                 Flood           - All-members
  F2                00:00:05:00:00:03 Learn           0 ge-0/0/44.0
  F2                00:19:e2:50:7d:e0 Static          - Router
  Linux             *                 Flood           - All-members
  Linux             00:19:e2:50:7d:e0 Static          - Router
  Linux             00:30:48:90:54:89 Learn           0 ge-0/0/47.0
  T1                *                 Flood           - All-members
  T1                00:00:05:00:00:01 Learn           0 ge-0/0/46.0
  T1                00:00:5e:00:01:00 Static          - Router
  T1                00:19:e2:50:63:e0 Learn           0 ge-0/0/46.0
  T1                00:19:e2:50:7d:e0 Static          - Router
  T10               *                 Flood           - All-members
  T10               00:00:5e:00:01:09 Static          - Router
  T10               00:19:e2:50:63:e0 Learn           0 ge-0/0/46.0
  T10               00:19:e2:50:7d:e0 Static          - Router
  T111              *                 Flood           - All-members
  T111              00:19:e2:50:63:e0 Learn           0 ge-0/0/15.0
  T111              00:19:e2:50:7d:e0 Static          - Router
  T111              00:19:e2:50:ac:00 Learn           0 ge-0/0/15.0
  T2                *                 Flood           - All-members
  T2                00:00:5e:00:01:01 Static          - Router
  T2                00:19:e2:50:63:e0 Learn           0 ge-0/0/46.0
  T2                00:19:e2:50:7d:e0 Static          - Router
  T3                *                 Flood           - All-members
  T3                00:00:5e:00:01:02 Static          - Router
  T3                00:19:e2:50:63:e0 Learn           0 ge-0/0/46.0
  T3                00:19:e2:50:7d:e0 Static          - Router
  T4                *                 Flood           - All-members
  T4                00:00:5e:00:01:03 Static          - Router
  T4                00:19:e2:50:63:e0 Learn           0 ge-0/0/46.0
[output truncated]
```

show ethernet-switching table detail

```
user@switch> show ethernet-switching table detail
Ethernet-switching table: 5 entries, 2 learned
 VLAN: default, Tag: 0, MAC: *, Interface: All-members
  Interfaces:
   ge-0/0/11.0, ge-0/0/20.0, ge-0/0/30.0, ge-0/0/36.0, ge-0/0/3.0
  Type: Flood
  Nexthop index: 1307

 VLAN: default, Tag: 0, MAC: 00:1f:12:30:b8:83, Interface: ge-0/0/3.0
   Type: Learn, Age: 0, Learned: 20:09:26
   Nexthop index: 1315

 VLAN: v1, Tag: 101, MAC: *, Interface: All-members
  Interfaces:
   ge-0/0/31.0
  Type: Flood
  Nexthop index: 1313

 VLAN: v1, Tag: 101, MAC: 00:1f:12:30:b8:89, Interface: ge-0/0/31.0
  Type: Learn, Age: 0, Learned: 20:09:25
  Nexthop index: 1312
```

```
 VLAN: v2, Tag: 102, MAC: *, Interface: All-members
  Interfaces:
   ae0.0
  Type: Flood
  Nexthop index: 1317
```

show ethernet-switching table extensive

```
user@switch> show ethernet-switching table extensive
Ethernet-switching table: 3 entries, 1 learned

    VLAN: v1, Tag: 10, MAC: *, Interface: All-members
    Interfaces:
                ge-0/0/14.0, ge-0/0/1.0, ge-0/0/2.0, ge-0/0/3.0, ge-0/0/4.0,
                ge-0/0/5.0, ge-0/0/6.0, ge-0/0/7.0, ge-0/0/8.0, ge-0/0/10.0,
                ge-0/0/0.0
    Type: Flood
    Nexthop index: 567

    VLAN: v1, Tag: 10, MAC: 00:21:59:c6:93:22, Interface: Router
    Type: Static
    Nexthop index: 0

    VLAN: v1, Tag: 10, MAC: 00:21:59:c9:9a:4e, Interface: ge-0/0/14.0
    Type: Learn, Age: 0, Learned: 18:40:50
    Nexthop index: 564
```

show ethernet-switching table interface ge-0/0/1

```
user@switch> show ethernet-switching table interface ge-0/0/1
Ethernet-switching table: 1 unicast entries
  VLAN              MAC address       Type         Age Interfaces
  V1                *                 Flood          - All-members
  V1                00:00:05:00:00:05 Learn          0 ge-0/0/1.0
```

PART 19

Port Security

CHAPTER 93

Port Security Overview

Port Security for J-EX Series Switches Overview

Ethernet LANs are vulnerable to attacks such as address spoofing (forging) and Layer 2 denial of service (DoS) on network devices. Port security features help protect the access ports on your switch against the losses of information and productivity that can result from such attacks.

The Junos OS on J-EX Series Switches provides features to help secure ports on the switch. The ports can be categorized as either trusted or untrusted. You apply policies appropriate to those categories to protect against various types of attacks.

Port security features can be turned on to obtain the most robust port security level. Basic port security features are enabled in the switch's default configuration. You can configure additional features with minimal configuration steps.

Port security features on J-EX Series switches are:

- DHCP snooping—Filters and blocks ingress DHCP server messages on untrusted ports; builds and maintains an IP-address/MAC-address binding database (called the DHCP snooping database). You enable this feature on VLANs.
- Dynamic ARP inspection (DAI)—Prevents ARP spoofing attacks. ARP requests and replies are compared against entries in the DHCP snooping database, and filtering

decisions are made based on the results of those comparisons. You enable this feature on VLANs.

- MAC limiting—Protects against flooding of the Ethernet switching table (also known as the MAC forwarding table or Layer 2 forwarding table). You enable this feature on access interfaces (ports).
- MAC move limiting—Detects MAC movement and MAC spoofing on access ports. You enable this feature on VLANs.
- Trusted DHCP server—With a DHCP server on a trusted port, protects against rogue DHCP servers sending leases. You enable this feature on interfaces (ports). By default, access ports are untrusted and trunk ports are trusted. (Access ports are the switch ports that connect to Ethernet endpoints such as user PCs and laptops, servers, and printers. Trunk ports are the switch ports that connect to other Ethernet switches or to routers.)
- IP source guard—Mitigates the effects of IP address spoofing attacks on the Ethernet LAN. You enable this feature on VLANs. With IP source guard enabled, the source IP address in the packet sent from an untrusted access interface is validated against the source MAC address in the DHCP snooping database. The packet is allowed for further processing if the source IP address to source MAC address binding is valid; if the binding is not valid, the packet is discarded.
- DHCP option 82—Also known as the DHCP relay agent information option. Helps protect the J-EX Series switch against attacks such as spoofing of IP addresses and MAC addresses and DHCP IP address starvation. Option 82 provides information about the network location of a DHCP client, and the DHCP server uses this information to implement IP addresses or other parameters for the client.

Related Documentation

- Security Features for J-EX Series Switches Overview on page 16
- Understanding DHCP Snooping for Port Security on J-EX Series Switches on page 2549
- Understanding DAI for Port Security on J-EX Series Switches on page 2555
- Understanding MAC Limiting and MAC Move Limiting for Port Security on J-EX Series Switches on page 2557
- Understanding IP Source Guard for Port Security on J-EX Series Switches on page 2563
- Understanding DHCP Option 82 for Port Security on J-EX Series Switches on page 2560
- Understanding How to Protect Access Ports on J-EX Series Switches from Common Attacks on page 2546

Understanding How to Protect Access Ports on J-EX Series Switches from Common Attacks

Port security features can protect the J-EX Series Switch against various types of attacks. Protection methods against some common attacks are:

- Mitigation of Ethernet Switching Table Overflow Attacks on page 2547
- Mitigation of Rogue DHCP Server Attacks on page 2547

Mitigation of Ethernet Switching Table Overflow Attacks

In an overflow attack on the Ethernet switching table, an intruder sends so many requests from new MAC addresses that the table cannot learn all the addresses. When the switch can no longer use information in the table to forward traffic, it is forced to broadcast messages. Traffic flow on the switch is disrupted, and packets are sent to all hosts on the network. In addition to overloading the network with traffic, the attacker might also be able to sniff that broadcast traffic.

To mitigate such attacks, configure both a MAC limit for learned MAC addresses and some specific allowed MAC addresses. Use the MAC limit feature to control the total number of MAC addresses that can be added to the Ethernet switching table for the specified interface or interfaces. By setting the MAC addresses that are explicitly allowed, you ensure that the addresses of network devices whose network access is critical are guaranteed to be included in the Ethernet switching table. See “Example: Configuring MAC Limiting, Including Dynamic and Allowed MAC Addresses, to Protect the Switch from Ethernet Switching Table Overflow Attacks” on page 2576.

Mitigation of Rogue DHCP Server Attacks

If an attacker sets up a rogue DHCP server to impersonate a legitimate DHCP server on the LAN, the rogue server can start issuing leases to the network's DHCP clients. The information provided to the clients by this rogue server can disrupt their network access, causing DoS. The rogue server might also assign itself as the default gateway device for the network. The attacker can then sniff the network traffic and perpetrate a man-in-the-middle attack—that is, it misdirects traffic intended for a legitimate network device to a device of its choice.

To mitigate a rogue DHCP server attack, set the interface to which that rogue server is connected as untrusted. That action will block all ingress DHCP server messages from that interface. See “Example: Configuring a DHCP Server Interface as Untrusted to Protect the Switch from Rogue DHCP Server Attacks” on page 2579.

NOTE: The switch logs all DHCP server packets that are received on untrusted ports—for example:

```
5 untrusted DHCPOFFER received, interface ge-0/0/0.0[65], vlan v1[10] server
ip/mac 12.12.12.1/00:00:00:00:01:12 offer ip/client mac
12.12.12.253/00:AA:BB:CC:DD:01
```

You can use these messages to detect malicious DHCP servers on the network.

Protection Against ARP Spoofing Attacks

In ARP spoofing, an attacker sends faked ARP messages on the network. The attacker associates its own MAC address with the IP address of a network device connected to the switch. Any traffic sent to that IP address is instead sent to the attacker. Now the attacker can create various types of mischief, including sniffing the packets that were meant for another host and perpetrating man-in-the middle attacks. (In a man-in-the-middle attack, the attacker intercepts messages between two hosts, reads them, and perhaps alters them, all without the original hosts knowing that their communications have been compromised.)

To protect against ARP spoofing on your switch, enable both DHCP snooping and dynamic ARP inspection (DAI). DHCP snooping builds and maintains the DHCP snooping table. That table contains the MAC addresses, IP addresses, lease times, binding types, VLAN information, and interface information for the untrusted interfaces on the switch. DAI uses the information in the DHCP snooping table to validate ARP packets. Invalid ARP packets are blocked and, when they are blocked, a system log message is recorded that includes the type of ARP packet and the sender's IP address and MAC address.

See "Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks" on page 2586.

Protection Against DHCP Snooping Database Alteration Attacks

In an attack designed to alter the DHCP snooping database, an intruder introduces a DHCP client on one of the switch's untrusted access interfaces that has a MAC address identical to that of a client on another untrusted port. The intruder acquires the DHCP lease, which results in changes to the entries in the DHCP snooping table. Subsequently, what would have been valid ARP requests from the legitimate client are blocked.

To protect against this type of alteration of the DHCP snooping database, configure MAC addresses that are explicitly allowed on the interface. See "Example: Configuring Allowed MAC Addresses to Protect the Switch from DHCP Snooping Database Alteration Attacks" on page 2590.

Protection Against DHCP Starvation Attacks

In a DHCP starvation attack, an attacker floods an Ethernet LAN with DHCP requests from spoofed (counterfeit) MAC addresses so that the switch's trusted DHCP servers cannot keep up with requests from legitimate DHCP clients on the switch. The address space of those servers is completely used up, so they can no longer assign IP addresses and lease times to clients. DHCP requests from those clients are either dropped—that is, the result is a denial of service (DoS)—or directed to a rogue DHCP server set up by the attacker to impersonate a legitimate DHCP server on the LAN.

To protect the switch from DHCP starvation attacks, use the MAC limiting feature. Specify the maximum number of MAC addresses that the switch can learn on the access interfaces to which those clients connect. The switch's DHCP server or servers will then be able to supply the specified number of IP addresses and leases to those clients and no more. If a DHCP starvation attack occurs after the maximum number of IP addresses has been

assigned, the attack will fail. See “Example: Configuring MAC Limiting to Protect the Switch from DHCP Starvation Attacks” on page 2583.

Related Documentation

- Understanding DHCP Snooping for Port Security on J-EX Series Switches on page 2549
- Understanding DAI for Port Security on J-EX Series Switches on page 2555
- Understanding MAC Limiting and MAC Move Limiting for Port Security on J-EX Series Switches on page 2557
- Understanding Trusted DHCP Servers for Port Security on J-EX Series Switches on page 2559
- Configuring Port Security (CLI Procedure) on page 2626
- Configuring Port Security (J-Web Procedure) on page 2627

Understanding DHCP Snooping for Port Security on J-EX Series Switches

DHCP snooping allows the switch to monitor and control DHCP messages received from untrusted devices connected to the switch. When DHCP snooping is enabled, the system snoops the DHCP messages to view DHCP lease information and build and maintain a database of valid IP address to MAC address (IP-MAC) bindings called the DHCP snooping database. Only clients with valid bindings are allowed access to the network.

DHCP Snooping Basics

Dynamic Host Configuration Protocol (DHCP) allocates IP addresses dynamically, “leasing” addresses to devices so that the addresses can be reused when no longer needed. Hosts and end devices that require IP addresses obtained through DHCP must communicate with a DHCP server across the LAN.

DHCP snooping acts as a guardian of network security by keeping track of valid IP addresses assigned to downstream network devices by a trusted DHCP server (the server is connected to a trusted network port). By default, all trunk ports on the switch are trusted and all access ports are untrusted for DHCP snooping. You can modify these defaults on each of the switch's interfaces.

When DHCP snooping is enabled, the lease information from the switch (which is a DHCP client) is used to create the DHCP snooping database, a mapping of IP address to VLAN–MAC-address pairs. For each VLAN–MAC-address pair, the database stores the corresponding IP address.

Entries in the DHCP database are updated in these events:

- When a DHCP client releases an IP address (sends a DHCPRELEASE message), the associated mapping entry is deleted from the database.
- If you move a network device from one VLAN to another, typically the device has to acquire a new IP address, so its entry in the database, including the VLAN ID, is updated.
- When the lease time (timeout value) assigned by the DHCP server expires, the associated entry is deleted from the database.

TIP: By default, the IP-MAC bindings are lost when the switch is rebooted and DHCP clients (the network devices, or hosts) must reacquire bindings. However, you can configure the bindings to persist by setting the **dhcp-snooping-file** statement to store the database file either locally or remotely.

You can configure the switch to snoop DHCP server responses only from particular VLANs. Doing this prevents spoofing of DHCP server messages.

You configure DHCP snooping for each VLAN, not for each interface (port). By default, DHCP snooping is disabled for all VLANs.

TIP: For private VLANs (PVLANs), enable DHCP snooping on the primary VLAN. If you enable DHCP snooping only on a community VLAN, DHCP messages coming from PVLAN trunk ports are not snooped.

DHCP Snooping Process

The basic process of DHCP snooping is shown in Figure 60 on page 2551.

Figure 60: DHCP Snooping

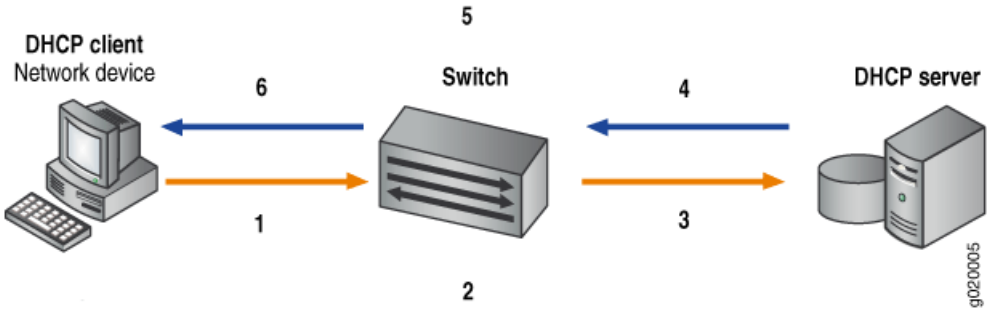

1. Device sends DHCPDISCOVER to request IP address or DHCPREQUEST to accept IP address and lease.
2. Switch snoops packet. Adds IP-MAC placeholder binding to database.
3. Switch forwards DHCPDISCOVER or DHCPREQUEST.
4. Server sends DHCPOFFER to offer address, DHCPACK to assign one, or DHCPNAK to deny address request.
5. Switch snoops packet. If placeholder exists, replaces it with IP-MAC binding on receipt of DHCPACK.
6. Switch forwards DHCPOFFER, DHCPACK, or DHCPNAK.

.

For general information about the messages that the DHCP client and DHCP server exchange during the assignment of an IP address for the client, see the *Junos OS System Basics Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.

DHCP Server Access

Switch access to the DHCP server can be configured in three ways:

- Switch, DHCP Clients, and DHCP Server Are All on the Same VLAN on page 2551
- Switch Acts as DHCP Server on page 2553
- Switch Acts as Relay Agent on page 2553

Switch, DHCP Clients, and DHCP Server Are All on the Same VLAN

When the switch, DHCP clients, and DHCP server are all members of the same VLAN, the DHCP server can be connected to the switch in one of two ways:

- The server is directly connected to the same switch as the one connected to the DHCP clients (the hosts, or network devices, that are requesting IP addresses from the server). You must configure the port that connects the server to the switch as a trusted port. See Figure 61 on page 2552.
- The server is directly connected to a switch that is itself directly connected through a trunk port to the switch that the DHCP clients are connected to. The trunk port is configured by default as a trusted port. The switch that the DHCP server is connected

to is not configured for DHCP snooping. See Figure 62 on page 2552—in the figure, **ge-0/0/11** is a trusted trunk port.

Figure 61: DHCP Server Connected Directly to Switch

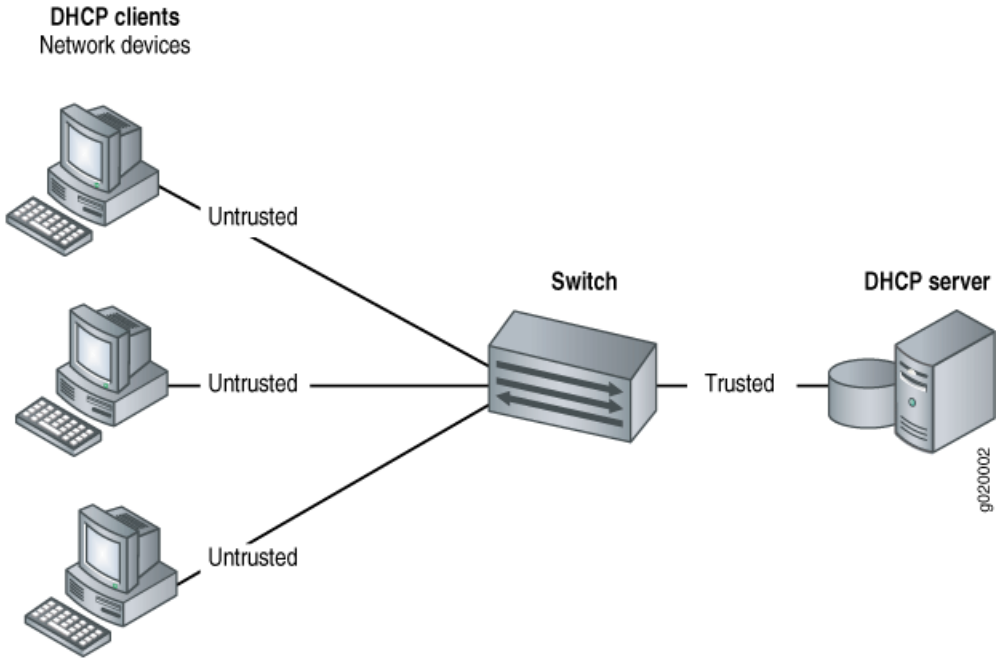

Figure 62: DHCP Server Connected Directly to Switch 2, with Switch 2 Connected to Switch 1 Through a Trusted Trunk Port

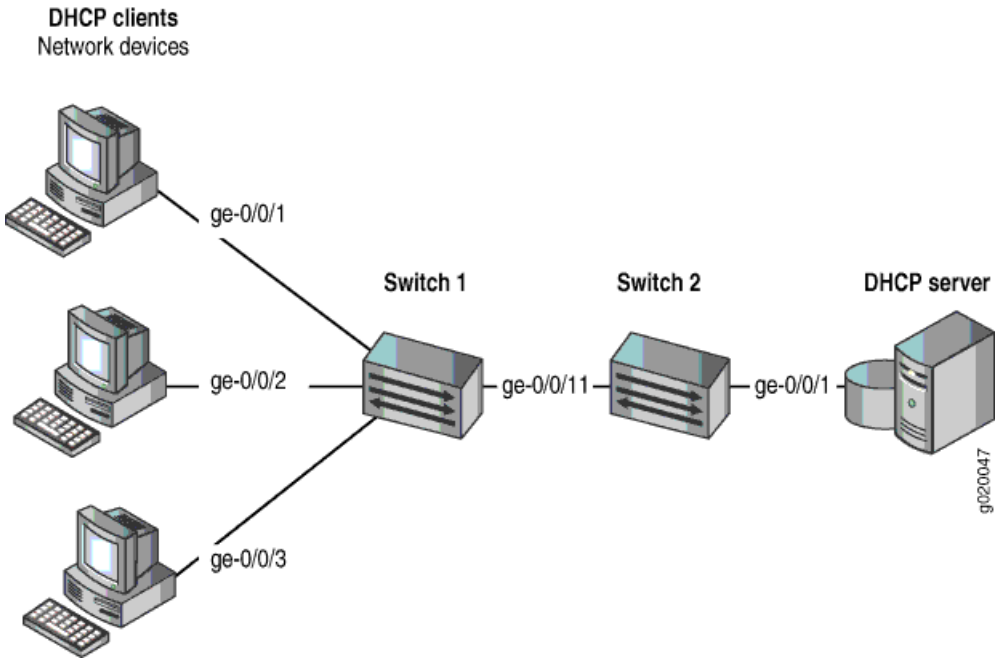

Switch Acts as DHCP Server

The switch itself is configured as a DHCP server; this is known as a “local” configuration. See Figure 63 on page 2553.

Figure 63: Switch Is the DHCP Server

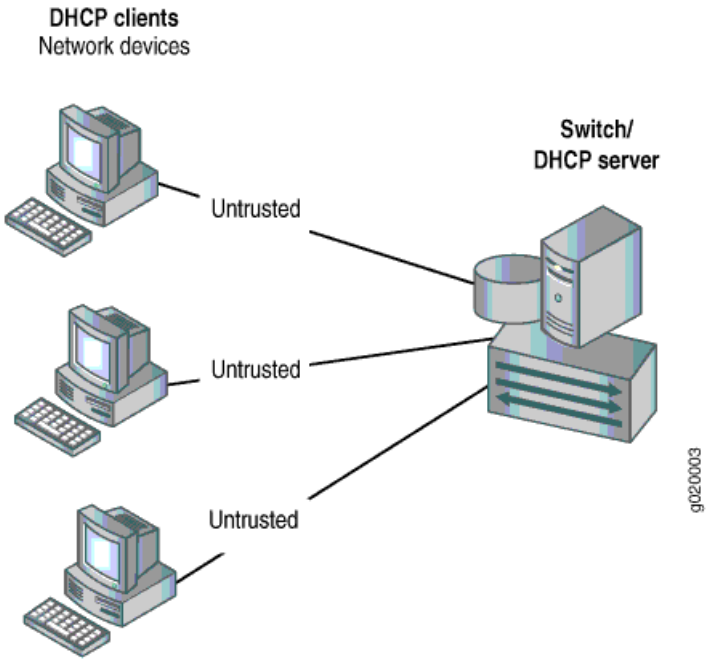

Switch Acts as Relay Agent

The switch functions as a relay agent when the DHCP clients or the DHCP server is connected to the switch through a Layer 3 interface (on the switch, these interfaces are configured as routed VLAN interfaces, or RVIs). These trunk interfaces are trusted by default.

These two scenarios illustrate the switch acting as a relay agent:

- The DHCP server and clients are in different VLANs.
- The switch is connected to a router that is in turn connected to the DHCP server. See Figure 64 on page 2554.

Figure 64: Switch Acting as Relay Agent Through Router to DHCP Server

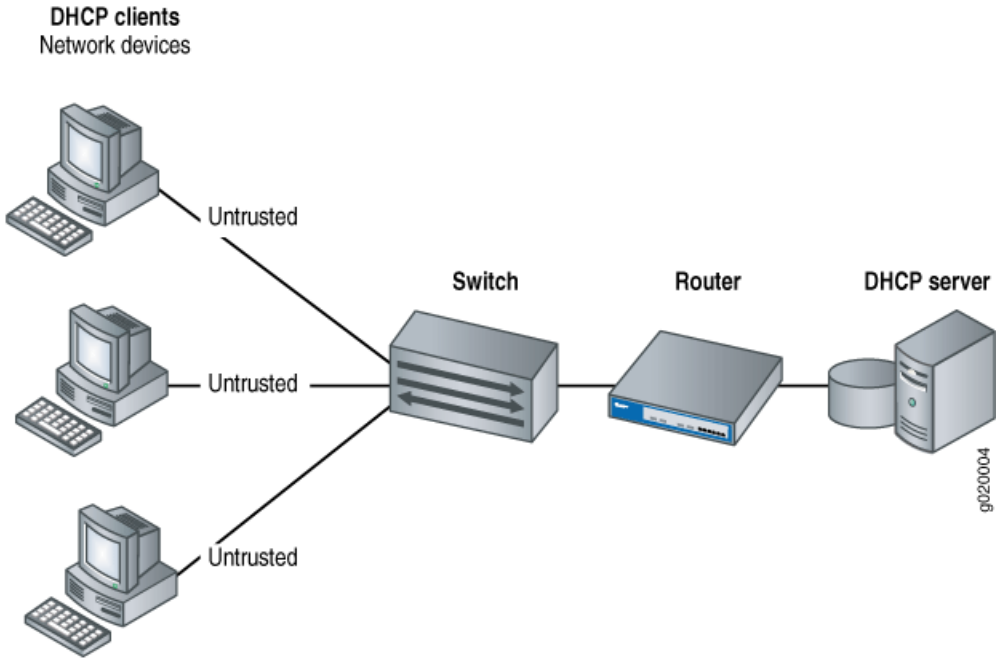

DHCP Snooping Table

The software creates a DHCP snooping information table that displays the content of the DHCP snooping database. The table shows current IP-MAC bindings, as well as lease time, type of binding, names of associated VLANs, and associated interface. To view the table, type **show dhcp snooping binding** at the operational mode prompt:

```
user@switch> show dhcp snooping binding
DHCP Snooping Information:
MAC address        IP address   Lease (seconds)  Type     VLAN      Interface
00:05:85:3A:82:77  192.0.2.17   600              dynamic  employee  ge-0/0/1.0
00:05:85:3A:82:79  192.0.2.18   653              dynamic  employee  ge-0/0/1.0
00:05:85:3A:82:80  192.0.2.19   720              dynamic  employee  ge-0/0/2.0
```

Static IP Address Additions to the DHCP Snooping Database

You can add specific static IP addresses to the database as well as have the addresses dynamically assigned through DHCP snooping. To add static IP addresses, you supply the IP address, the MAC address of the device, the interface on which the device is connected, and the VLAN with which the interface is associated. No lease time is assigned to the entry. The statically configured entry never expires.

Snooping DHCP Packets That Have Invalid IP Addresses

If you enable DHCP snooping on a VLAN and then devices on that VLAN send DHCP packets that request invalid IP addresses, these invalid IP addresses will be stored in the DHCP snooping database until they are deleted when their default timeout is reached. To eliminate this unnecessary consumption of space in the DHCP snooping database,

the switch drops the DCHP packets that request invalid IP addresses, preventing the snooping of these packets. The invalid IP addresses are:

- 0.0.0.0
- 128.0.x.x
- 191.255.x.x
- 192.0.0.x
- 223.255.255.x
- 224.x.x.x
- 240.x.x.x to 255.255.255.255

Related Documentation

- Port Security for J-EX Series Switches Overview on page 2545
- Understanding Trusted DHCP Servers for Port Security on J-EX Series Switches on page 2559
- Understanding DHCP Option 82 for Port Security on J-EX Series Switches on page 2560
- DHCP Services for J-EX Series Switches Overview on page 445
- DHCP/BOOTP Relay for J-EX Series Switches Overview on page 446
- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Enabling DHCP Snooping (CLI Procedure) on page 2630 and Enabling DHCP Snooping (J-Web Procedure) on page 2631
- Troubleshooting Port Security on page 2665

Understanding DAI for Port Security on J-EX Series Switches

Dynamic ARP inspection (DAI) protects J-EX Series Switches against ARP spoofing.

DAI inspects ARP packets on the LAN and uses the information in the DHCP snooping database on the switch to validate ARP packets and to protect against ARP cache poisoning. ARP requests and replies are compared against entries in the DHCP snooping database, and filtering decisions are made based on the results of those comparisons. When an attacker tries to use a forged ARP packet to spoof an address, the switch compares the address to entries in the database. If the MAC address or IP address in an ARP packet does not match a valid entry in the DHCP snooping database, the packet is dropped.

ARP packets are trapped to the Routing Engine and are rate-limited to protect the switch from CPU overload.

- Address Resolution Protocol on page 2556
- ARP Spoofing on page 2556
- DAI on J-EX Series Switches on page 2556

Address Resolution Protocol

Sending IP packets on a multiaccess network requires mapping an IP address to an Ethernet media access control (MAC) address.

Ethernet LANs use Address Resolution Protocol (ARP) to map MAC addresses to IP addresses.

The switch maintains this mapping in a cache that it consults when forwarding packets to network devices. If the ARP cache does not contain an entry for the destination device, the host (the DHCP client) broadcasts an ARP request for that device's address and stores the response in the cache.

ARP Spoofing

ARP spoofing (also known as ARP poisoning or ARP cache poisoning) is one way to initiate man-in-the-middle attacks. The attacker sends an ARP packet that spoofs the MAC address of another device on the LAN. Instead of the switch sending traffic to the proper network device, it sends it to the device with the spoofed address that is impersonating the proper device. If the impersonating device is the attacker's machine, the attacker receives all the traffic from the switch that should have gone to another device. The result is that traffic from the switch is misdirected and cannot reach its proper destination.

One type of ARP spoofing is gratuitous ARP, which is when a network device sends an ARP request to resolve its own IP address. In normal LAN operation, gratuitous ARP messages indicate that two devices have the same MAC address. They are also broadcast when a network interface card (NIC) in a device is changed and the device is rebooted, so that other devices on the LAN update their ARP caches. In malicious situations, an attacker can poison the ARP cache of a network device by sending an ARP response to the device that directs all packets destined for a certain IP address to go to a different MAC address instead.

To prevent MAC spoofing through gratuitous ARP and through other types of spoofing, J-EX Series switches examine ARP responses through DAI.

DAI on J-EX Series Switches

DAI examines ARP requests and responses on the LAN and validates ARP packets. The switch intercepts ARP packets from an access port and validates them against the DHCP snooping database. If no IP-MAC entry in the database corresponds to the information in the ARP packet, DAI drops the ARP packet and the local ARP cache is not updated with the information in that packet. DAI also drops ARP packets when the IP address in the packet is invalid.

The Junos OS for J-EX switches uses DAI for ARP packets received on access ports because these ports are untrusted by default. Trunk ports are trusted by default, so ARP packets bypass DAI on them.

You configure DAI for each VLAN, not for each interface (port). By default, DAI is disabled for all VLANs. You can set an interface to be trusted for ARP packets by setting **dhcp-trusted** on that port.

For packets directed to the switch to which a network device is connected, ARP queries are broadcast on the VLAN. The ARP responses to those queries are subjected to the DAI check.

For DAI, all ARP packets are trapped to the Routing Engine. To prevent CPU overloading, ARP packets destined for the Routing Engine are rate-limited.

If the DHCP server goes down and the lease time for an IP-MAC entry for a previously valid ARP packet runs out, that packet is blocked.

Related Documentation

- Port Security for J-EX Series Switches Overview on page 2545
- Understanding DHCP Snooping for Port Security on J-EX Series Switches on page 2549
- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Example: Configuring DHCP Snooping, DAI , and MAC Limiting on a J-EX Series Switch with Access to a DHCP Server Through a Second Switch on page 2593
- Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks on page 2586
- Enabling Dynamic ARP Inspection (CLI Procedure) on page 2633
- Enabling Dynamic ARP Inspection (J-Web Procedure) on page 2634

Understanding MAC Limiting and MAC Move Limiting for Port Security on J-EX Series Switches

MAC limiting protects against flooding of the Ethernet switching table (also known as the MAC forwarding table or Layer 2 forwarding table). You enable this feature on interfaces (ports). MAC move limiting detects MAC movement and MAC spoofing on access interfaces. You enable this feature on VLANs.

MAC Limiting

MAC limiting sets a limit on the number of MAC addresses that can be learned on a single Layer 2 access interface or on all the Layer 2 access interfaces on the switch. The Junos OS provides two MAC limiting methods:

- Maximum number of MAC addresses—You configure the maximum number of dynamic MAC addresses allowed per interface. When the limit is exceeded, incoming packets with new MAC addresses are treated as specified by the configuration. The incoming packets with new MAC addresses can be ignored, dropped, logged, or the interface can be shut down or temporarily disabled.

- Allowed MAC—You configure specific “allowed” MAC addresses for the access interface. Any MAC address that is not in the list of configured addresses is not learned and the switch logs the message. Allowed MAC binds MAC addresses to a VLAN so that the address does not get registered outside the VLAN. If an allowed MAC setting conflicts with a dynamic MAC setting, the allowed MAC setting takes precedence.

NOTE: If you do not want the switch to log messages received for invalid MAC addresses on an interface that has been configured for specific “allowed” MAC addresses, you can disable the logging by configuring the no-allowed-mac-log statement.

MAC Move Limiting

MAC move limiting causes the switch to track the number of times a MAC address can move to a new interface (port). It can help to prevent MAC spoofing, and it can also detect and prevent loops.

If a MAC address moves more than the configured number of times within one second, the switch performs the configured action. You can configure MAC move limiting to apply to all VLANs or to a specific VLAN.

Actions for MAC Limiting and MAC Move Limiting

You can choose to have one of the following actions performed when the limit of MAC addresses or the limit of MAC moves is exceeded:

- **drop**—Drop the packet and generate an alarm, an SNMP trap, or a system log entry. This is the default.
- **log**—Do not drop the packet but generate an alarm, an SNMP trap, or a system log entry.
- **none**—Take no action.
- **shutdown**—Disable the interface and generate an alarm. If you have configured the switch with the **port-error-disable** statement, the disabled interface recovers automatically upon expiration of the specified disable timeout. If you have not configured the switch for autorecovery from port error disabled conditions, you can bring up the disabled interfaces by running the **clear ethernet-switching port-error** command.

See descriptions of results of these various action settings in “Verifying That MAC Limiting Is Working Correctly” on page 2657.

If you have set a MAC limit to apply to all interfaces on the switch, you can override that setting for a particular interface by specifying action **none**. See “Setting the none Action on an Interface to Override a MAC Limit Applied to All Interfaces (CLI Procedure)” on page 2642.

MAC Addresses That Exceed the MAC Limit or MAC Move Limit

If you have configured the **port-error-disable** statement, you can view which interfaces are temporarily disabled due to exceeding the MAC limit or MAC move limit in the output for the **show ethernet-switching interfaces** command.

The log messages that indicate the MAC limit or MAC move limit has been exceeded include the offending MAC addresses that have exceeded the limit. See "Troubleshooting Port Security" on page 2665 for details.

Related Documentation

- Port Security for J-EX Series Switches Overview on page 2545
- Example: Configuring MAC Limiting, Including Dynamic and Allowed MAC Addresses, to Protect the Switch from Ethernet Switching Table Overflow Attacks on page 2576
- Example: Configuring MAC Limiting to Protect the Switch from DHCP Starvation Attacks on page 2583
- Configuring MAC Limiting (CLI Procedure) on page 2635
- Configuring MAC Limiting (J-Web Procedure) on page 2637
- Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure) on page 2516
- **no-allowed-mac-log on page 2686**

Understanding Trusted DHCP Servers for Port Security on J-EX Series Switches

Any interface on the switch that connects to a DHCP server can be configured as a trusted port. Configuring a DHCP server on a trusted port protects against rogue DHCP servers sending leases.

Ensure that the DHCP server interface is physically secure—that is, that access to the server is monitored and controlled at the site—before you configure the port as trusted.

Related Documentation

- Understanding DHCP Snooping for Port Security on J-EX Series Switches on page 2549
- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Example: Configuring a DHCP Server Interface as Untrusted to Protect the Switch from Rogue DHCP Server Attacks on page 2579
- Enabling a Trusted DHCP Server (CLI Procedure) on page 2632
- Enabling a Trusted DHCP Server (J-Web Procedure) on page 2632

Understanding DHCP Option 82 for Port Security on J-EX Series Switches

You can use DHCP option 82, also known as the DHCP relay agent information option, to help protect the switch against attacks such as spoofing (forging) of IP addresses and MAC addresses, and DHCP IP address starvation. Hosts on untrusted access interfaces on Ethernet LAN switches send requests for IP addresses in order to access the Internet. The switch forwards or relays these requests to DHCP servers, and the servers send offers for IP address leases in response. Attackers can use these messages to perpetrate address spoofing and starvation.

Option 82 provides information about the network location of a DHCP client, and the DHCP server uses this information to implement IP addresses or other parameters for the client. The Junos OS implementation of DHCP option 82 supports RFC 3046, *DHCP Relay Agent Information Option*, at http://tools.ietf.org/html/rfc3046.

This topic covers:

DHCP Option 82 Processing

If DHCP option 82 is enabled on the switch, then when a network device—a DHCP client—that is connected to the switch on an untrusted interface sends a DHCP request, the switch inserts information about the client's network location into the packet header of that request. The switch then sends the request to the DHCP server. The DHCP server reads the option 82 information in the packet header and uses it to implement the IP address or another parameter for the client. See “Suboption Components of Option 82” on page 2561 for details about option 82 information.

You can enable DHCP option 82 on a single VLAN or on all VLANs on the switch. You can also configure it on Layer 3 interfaces (in routed VLAN interfaces, or RVIs) when the switch is functioning as a relay agent.

When option 82 is enabled on the switch, then this sequence of events occurs when a DHCP client sends a DHCP request:

1. The switch receives the request and inserts the option 82 information in the packet header.
2. The switch forwards or relays the request to the DHCP server.
3. The server uses the DHCP option 82 information to formulate its reply and sends a response back to the switch. It does not alter the option 82 information.
4. The switch strips the option 82 information from the response packet.
5. The switch forwards the response packet to the client.

NOTE: To use the DHCP option 82 feature, you must ensure that the DHCP server is configured to accept option 82. If it is not configured to accept option 82, then when it receives requests containing option 82 information, it does not use the information in setting parameters and it does not echo the information in its response message. For detailed information about configuring DHCP services, see the *Junos OS System Basics Configuration Guide* at http://www.juniper.net/techpubs/software/junos/. The configuration for DHCP service on the J-EX Series Switch includes the **dhcp** statement at the **[edit system services]** hierarchy level.

Suboption Components of Option 82

Option 82 as implemented on the J-EX Series switch comprises the suboptions circuit ID, remote ID, and vendor ID. These suboptions are fields in the packet header:

- circuit ID—Identifies the circuit (interface and/or VLAN) on the switch on which the request was received. The circuit ID contains the interface name and/or VLAN name, with the two elements separated by a colon—for example, **ge-0/0/10:vlan1**, where **ge-0/0/10** is the interface name and **vlan1** is the VLAN name. If the request packet is received on a Layer 3 interface, the circuit ID is just the interface name—for example, **ge-0/0/10**.

  Use the **prefix** option to add an optional prefix to the circuit ID. If you enable the **prefix** option, the hostname for the switch is used as the prefix; for example, **switch1:ge-0/0/10:vlan1**, where **switch1** is the hostname.

  You can also specify that the interface description be used rather than the interface name and/or that the VLAN ID be used rather than the VLAN name.

- remote ID—Identifies the host. By default, the remote ID is the MAC address of the switch. You can specify that the remote ID be the hostname of the switch, the interface description, or a character string of your choice. You can also add an optional prefix to the remote ID.

- vendor ID—Identifies the vendor of the host. If you specify the **vendor-id** option but do not enter a value, the default value **Juniper** is used. To specify a value, you type a character string.

Configurations of the J-EX Series Switch That Support Option 82

Configurations of the J-EX Series switch that support option 82 are:

- Switch and Clients Are on Same VLAN as DHCP Server on page 2561
- Switch Acts as Relay Agent on page 2562

Switch and Clients Are on Same VLAN as DHCP Server

If the DHCP clients, the switch, and the DHCP server are all on the same VLAN, the switch forwards the requests from the clients on untrusted access interfaces to the server on a trusted interface. See Figure 65 on page 2562.

Figure 65: DHCP Clients, Switch, and DHCP Server Are All on Same VLAN

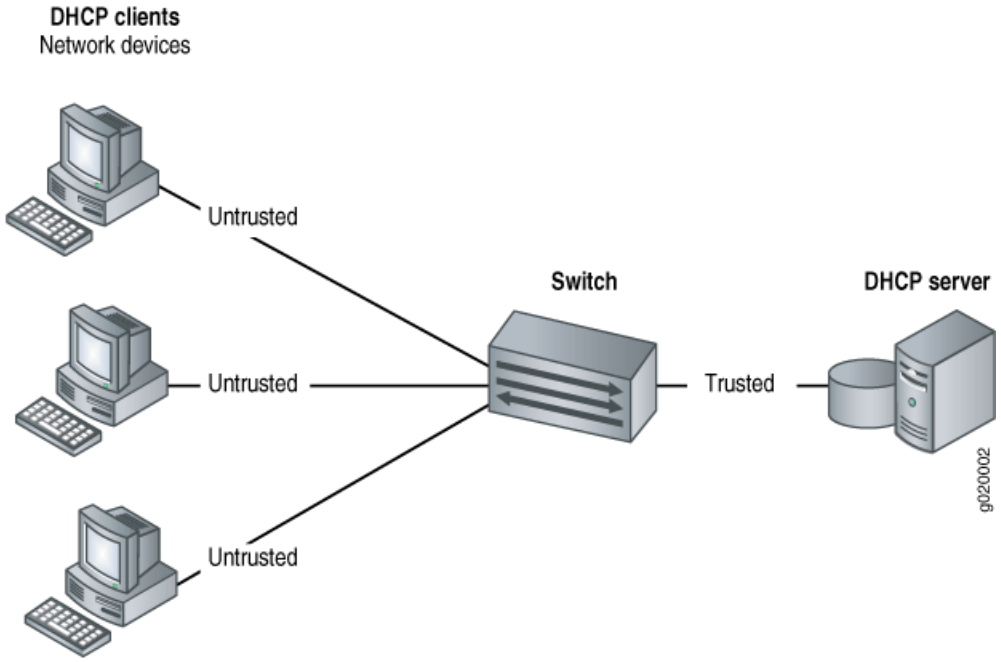

For the configuration shown in Figure 65 on page 2562, you set DHCP option 82 at the **[edit ethernet-switching-options secure-access-port vlan]** hierarchy level.

Switch Acts as Relay Agent

The switch functions as a relay agent when the DHCP clients or the DHCP server is connected to the switch through a Layer 3 interface. On the switch, these interfaces are configured as routed VLAN interfaces, or RVIs. Figure 66 on page 2563 illustrates a scenario for the switch-as-relay-agent; in this instance, the switch relays requests through a router to the server.

Figure 66: Switch Relays DHCP Requests to Server

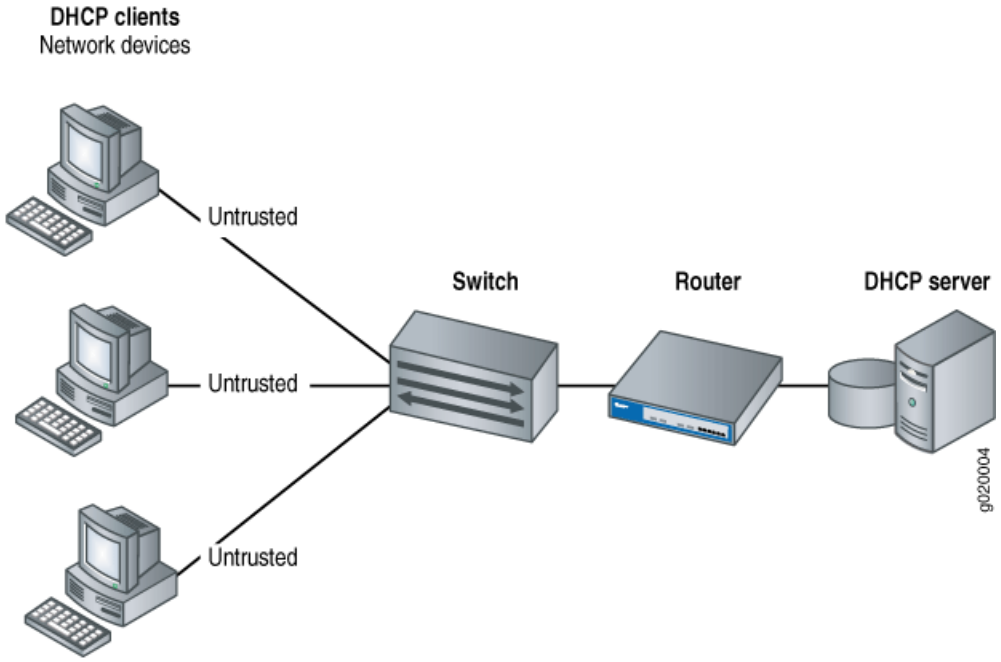

For the configuration shown in Figure 66 on page 2563, you set DHCP option 82 at the **[edit forwarding-options helpers bootp]** hierarchy level.

Related Documentation

- Port Security for J-EX Series Switches Overview on page 2545
- Example: Setting Up DHCP Option 82 on a J-EX Series Switch with No Relay Agent Between Clients and DHCP Server on page 2617
- Example: Setting Up DHCP Option 82 with a J-EX Series Switch as Relay Agent Between Clients and a DHCP Server on page 2615
- Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 2649
- Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 2646

Understanding IP Source Guard for Port Security on J-EX Series Switches

Ethernet LAN switches are vulnerable to attacks that involve spoofing (forging) of source IP addresses or source MAC addresses. You can use the IP source guard access port security feature on J-EX Series Switches to mitigate the effects of these attacks.

- IP Address Spoofing on page 2564
- How IP Source Guard Works on page 2564
- The IP Source Guard Database on page 2564
- Typical Uses of Other Junos OS Features with IP Source Guard on page 2565

IP Address Spoofing

Hosts on access interfaces can spoof source IP addresses and/or source MAC addresses by flooding the switch with packets containing invalid addresses. Such attacks combined with other techniques such as TCP SYN flood attacks can result in denial-of-service (DoS) attacks. With source IP address or source MAC address spoofing, the system administrator cannot identify the source of the attack. The attacker can spoof addresses on the same subnet or on a different subnet.

How IP Source Guard Works

IP source guard checks the IP source address and MAC source address in a packet sent from a host attached to an untrusted access interface on the switch against entries stored in the DHCP snooping database. If IP source guard determines that the packet header contains an invalid source IP address or source MAC address, it ensures that the switch does not forward the packet—that is, the packet is discarded.

When you configure IP source guard, you enable on it on one or more VLANs. IP source guard applies its checking rules to packets sent from untrusted access interfaces on those VLANs. By default, on J-EX Series switches, access interfaces are untrusted and trunk interfaces are trusted. IP source guard does not check packets that have been sent to the switch by devices connected to either trunk interfaces or trusted access interfaces—that is, interfaces configured as **dhcp-trusted** so that a DHCP server can be connected to that interface to provide dynamic IP addresses.

IP source guard obtains information about IP-address/MAC-address/VLAN bindings from the DHCP snooping database. It causes the switch to validate incoming IP packets against the entries in that database.

After the DHCP snooping database has been populated either through dynamic DHCP snooping or through configuration of specific static IP address/MAC address bindings, the IP source guard feature builds its database. It then checks incoming packets from access interfaces on the VLANs on which it is enabled. If the source IP addresses and source MAC addresses match the IP source guard binding entries, the switch forwards the packets to their specified destination addresses. If there are no matches, the switch discards the packets.

The IP Source Guard Database

The IP source guard database looks like this:

```
user@switch> show ip-source-guard
IP source guard information:
Interface    Tag  IP Address   MAC Address        VLAN

ge-0/0/12.0  0    10.10.10.7   00:30:48:92:A5:9D  vlan100

ge-0/0/13.0  0    10.10.10.9   00:30:48:8D:01:3D  vlan100

ge-0/0/13.0  100  *            *                  voice
```

The IP source guard database table contains the VLANs enabled for IP source guard, the untrusted access interfaces on those VLANs, the VLAN 802.1Q tag IDs if there are any,

and the IP addresses and MAC addresses that are bound to one another. If a switch interface is associated with multiple VLANs and some of those VLANs are enabled for IP source guard and others are not, the VLANs that are not enabled for IP source guard have a star (*) in the **IP Address** and **MAC Address** fields. See the entry for the **voice** VLAN in the preceding sample output.

Typical Uses of Other Junos OS Features with IP Source Guard

You can configure IP source guard with various other features on the J-EX Series switch to provide access port security, including:

- VLAN tagging (used for voice VLANs)
- GRES (Graceful Routing Engine switchover)
- Virtual Chassis configurations (multiple J-EX4200 switches that are managed through a single management interface)
- Link-aggregation groups (LAGs)
- 802.1X user authentication, in single supplicant mode

NOTE: The 802.1X user authentication is applied in one of three modes: single supplicant, single-secure supplicant, or multiple supplicant. Single supplicant mode works with IP source guard, but single-secure and multiple supplicant modes do not.

Related Documentation

- Understanding DHCP Snooping for Port Security on J-EX Series Switches on page 2549
- Example: Configuring IP Source Guard on a Data VLAN That Shares an Interface with a Voice VLAN on page 2608
- Example: Configuring IP Source Guard with Other J-EX Series Switch Features to Mitigate Address-Spoofing Attacks on Untrusted Access Interfaces on page 2600

Understanding Proxy ARP on J-EX Series Switches

You can configure proxy Address Resolution Protocol (ARP) on your J-EX Series Switch to enable the switch to respond to ARP queries for network addresses by offering its own Ethernet media access control (MAC) address. With proxy ARP enabled, the switch captures and routes traffic to the intended destination.

Proxy ARP is useful in situations where hosts are on different physical networks and you do not want to use subnet masking. Because ARP broadcasts are not propagated between hosts on different physical networks, hosts will not receive a response to their ARP request if the destination is on a different subnet. Enabling the switch to act as an ARP proxy allows the hosts to transparently communicate with each other through the switch. Proxy ARP can help hosts on a subnet reach remote subnets without your having to configure routing or a default gateway.

What Is ARP?

Ethernet LANs use ARP to map Ethernet MAC addresses to IP addresses. Each device maintains a cache containing a mapping of MAC addresses to IP addresses. The switch maintains this mapping in a cache that it consults when forwarding packets to network devices. If the ARP cache does not contain an entry for the destination device, the host (the DHCP client) broadcasts an ARP request for that device's address and stores the response in the cache.

Proxy ARP Overview

When proxy ARP is enabled, if the switch receives an ARP request for which it has a route to the target (destination) IP address, the switch responds by sending a proxy ARP reply packet containing its own MAC address. The host that sent the ARP request then sends its packets to the switch, which forwards them to the intended host.

NOTE: For security reasons, the source address in an ARP request must be on the same subnet as the interface on which the ARP request is received.

You can configure proxy ARP for each interface. You can also configure proxy ARP for a VLAN by using a routed VLAN interface (RVI).

J-EX Series switches support two modes of proxy ARP, restricted and unrestricted. Both modes require that the switch have an active route to the destination address of the ARP request.

- Restricted—The switch responds to ARP requests in which the physical networks of the source and target are different and does not respond if the source and target IP addresses are on the same subnet. In this mode, hosts on the same subnet communicate without proxy ARP. We recommend that you use this mode on the switch.

- Unrestricted—The switch responds to all ARP requests for which it has a route to the destination. This is the default mode (because it is the default mode in Junos OS configurations other than those on the switch). We recommend using restricted mode on the switch.

Best Practices for Proxy ARP on J-EX Series Switches

We recommend these best practices for configuring proxy ARP on the switches:

- Set proxy ARP to restricted mode.
- Use restricted mode when configuring proxy ARP on RVIs.
- If you set proxy ARP to unrestricted, disable gratuitous ARP requests on each interface enabled for proxy ARP.

Related Documentation

- Example: Configuring Proxy ARP on a J-EX Series Switch on page 2621
- Configuring Proxy ARP (CLI Procedure) on page 1153

CHAPTER 94

Examples: Port Security Configuration

Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch

You can configure DHCP snooping, dynamic ARP inspection (DAI), MAC limiting, and MAC move limiting on the access ports of J-EX Series switches to protect the switch and the Ethernet LAN against address spoofing and Layer 2 denial-of-service (DoS) attacks. You can also configure a trusted DHCP server and specific (allowed) MAC addresses for the switch interfaces.

This example describes how to configure basic port security features—DHCP snooping, DAI, MAC limiting, and MAC move limiting, as well as a trusted DHCP server and allowed MAC addresses—on a switch. The DHCP server and its clients are all members of a single VLAN on the switch.

Requirements

This example uses the following hardware and software components:

- One J-EX Series switch
- A DHCP server to provide IP addresses to network devices on the switch

Before you configure DHCP snooping, DAI, and MAC limiting port security features, be sure you have:

- Connected the DHCP server to the switch.
- Configured the VLAN **employee-vlan** on the switch. See “Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches” on page 1070.

Overview and Topology

Ethernet LANs are vulnerable to address spoofing and DoS attacks on network devices. To protect the devices from such attacks, you can configure DHCP snooping to validate DHCP server messages, DAI to protect against MAC spoofing, and MAC cache limiting to constrain the number of MAC addresses the switch adds to its MAC address cache. You can also configure MAC move limiting to help prevent MAC spoofing.

This example shows how to configure these security features on a J-EX4200-24T switch. The switch is connected to a DHCP server.

The setup for this example includes the VLAN **employee-vlan** on the switch. The procedure for creating that VLAN is described in the topic “Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches” on page 1070. That procedure is not repeated here. Figure 67 on page 2571 illustrates the topology for this example.

Figure 67: Network Topology for Basic Port Security

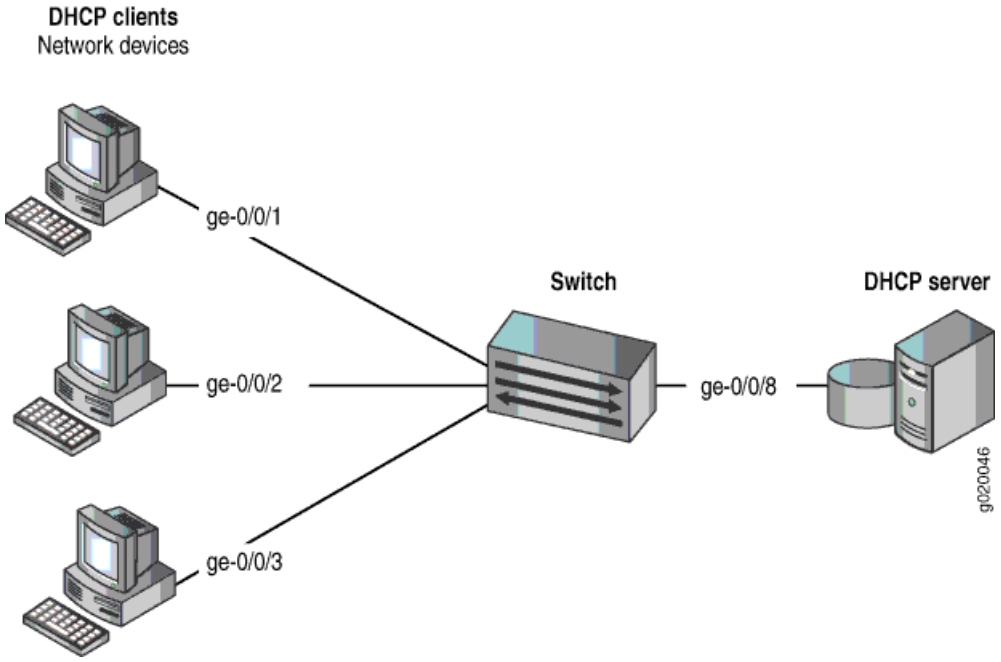

The components of the topology for this example are shown in Table 321 on page 2571.

Table 321: Components of the Port Security Topology

Properties	Settings
Switch hardware	One J-EX4200-24T, 24 ports (8 PoE ports)
VLAN name and ID	**employee-vlan**, tag **20**
VLAN subnets	**192.0.2.16/28** **192.0.2.17** through **192.0.2.30** **192.0.2.31** is subnet's broadcast address
Interfaces in **employee-vlan**	**ge-0/0/1**, **ge-0/0/2**, **ge-0/0/3**, **ge-0/0/8**
Interface for DHCP server	**ge-0/0/8**

In this example, the switch is initially configured with the default port security setup. In the default configuration on the switch:

- Secure port access is activated on the switch.
- DHCP snooping and DAI are disabled on all VLANs.
- All access ports are untrusted and all trunk ports are trusted for DHCP snooping, which is the default setting.

In the configuration tasks for this example, you set the DHCP server first as untrusted and then as trusted; you enable DHCP snooping, DAI, and MAC move limiting on a VLAN;

you modify the value for MAC limit; and you configure some specific (allowed) MAC addresses on an interface.

Configuration

To configure basic port security on a switch whose DHCP server and client ports are in a single VLAN:

CLI Quick Configuration

To quickly configure basic port security on the switch, copy the following commands and paste them into the switch terminal window:

```
[edit ethernet-switching-options secure-access-port]
set interface ge-0/0/1 mac-limit 4
set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:80
set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:81
set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:83
set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:85
set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:88
set interface ge-0/0/2 mac-limit 4
set interface ge-0/0/8 dhcp-trusted
set vlan employee–vlan arp-inspection
set vlan employee-vlan examine-dhcp
set vlan employee-vlan mac-move-limit 5
```

Step-by-Step Procedure

Configure basic port security on the switch:

1. Enable DHCP snooping on the VLAN:

   ```
   [edit ethernet-switching-options secure-access-port]
   user@switch# set vlan employee-vlan examine-dhcp
   ```

2. Specify the interface (port) from which DHCP responses are allowed:

   ```
   [edit ethernet-switching-options secure-access-port]
   user@switch# set interface ge-0/0/8 dhcp-trusted
   ```

3. Enable dynamic ARP inspection (DAI) on the VLAN:

   ```
   [edit ethernet-switching-options secure-access-port]
   user@switch# set vlan employee-vlan arp-inspection
   ```

4. Configure the MAC limit of **4** and use the default action, **drop**. (Packets will be dropped and the MAC address will not be added to the Ethernet switching table if the MAC limit has been exceeded on the interfaces):

   ```
   [edit ethernet-switching-options secure-access-port]
   user@switch# set interface ge-0/0/1 mac-limit 4
   user@switch# set interface ge-0/0/2 mac-limit 4
   ```

5. Configure a MAC move limit of **5** and use the default action, **drop**. (Packets will be dropped and the MAC address will not be added to the Ethernet switching table if a MAC address has exceeded the MAC move limit):

   ```
   [edit ethernet-switching-options secure-access-port]
   user@switch# set vlan employee-vlan mac-move-limit 5
   ```

6. Configure the allowed MAC addresses:

   ```
   [edit ethernet-switching-options secure-access-port]
   user@switch# set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:80
   user@switch# set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:81
   ```

```
user@switch# set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:83
user@switch# set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:85
user@switch# set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:88
```

Results Check the results of the configuration:

```
[edit ethernet-switching-options secure-access-port]
user@switch# show
interface ge-0/0/1.0 {
   mac-limit 4 action drop;
}
interface ge-0/0/2.0 {
   allowed-mac [ 00:05:85:3a:82:80 00:05:85:3a:82:81 00:05:85:3a:82:83
      00:05:85:3a:82:85 00:05:85:3a:82:88 ];
   mac-limit 4 action drop;
}
interface ge-0/0/8.0 {
   dhcp-trusted;
}
vlan employee-vlan {
   arp-inspection
   examine-dhcp;
   mac-move-limit 5 action drop;
}
```

Verification

To confirm that the configuration is working properly:

- Verifying That DHCP Snooping Is Working Correctly on the Switch on page 2573
- Verifying That DAI Is Working Correctly on the Switch on page 2574
- Verifying That MAC Limiting and MAC Move Limiting Are Working Correctly on the Switch on page 2574
- Verifying That Allowed MAC Addresses Are Working Correctly on the Switch on page 2575

Verifying That DHCP Snooping Is Working Correctly on the Switch

Purpose Verify that DHCP snooping is working on the switch.

Action Send some DHCP requests from network devices (here they are DHCP clients) connected to the switch.

Display the DHCP snooping information when the interface on which the DHCP server connects to the switch is trusted. The following output results when requests are sent from the MAC addresses and the server has provided the IP addresses and leases:

```
user@switch> show dhcp snooping binding
```

```
DHCP Snooping Information:
MAC Address         IP Address    Lease    Type     VLAN          Interface
-----------------   ----------    -----    ----     ----          ---------
00:05:85:3A:82:77    192.0.2.17    600    dynamic  employee-vlan   ge-0/0/1.0
00:05:85:3A:82:79    192.0.2.18    653    dynamic  employee-vlan   ge-0/0/1.0
00:05:85:3A:82:80    192.0.2.19    720    dynamic  employee-vlan   ge-0/0/2.0
00:05:85:3A:82:81    192.0.2.20    932    dynamic  employee-vlan   ge-0/0/2.0
00:05:85:3A:82:83    192.0.2.21   1230    dynamic  employee-vlan   ge-0/0/2.0
00:05:85:27:32:88    192.0.2.22   3200    dynamic  employee-vlan   ge-0/0/2.0
```

Meaning When the interface on which the DHCP server connects to the switch has been set to trusted, the output (see preceding sample) shows, for each MAC address, the assigned IP address and lease time—that is, the time, in seconds, remaining before the lease expires.

If the DHCP server had been configured as untrusted, no entries would be added to the DHCP snooping database and nothing would be shown in the output of the **show dhcp snooping binding** command.

Verifying That DAI Is Working Correctly on the Switch

Purpose Verify that DAI is working on the switch.

Action Send some ARP requests from network devices connected to the switch.

Display the DAI information:

```
user@switch> show arp inspection statistics
ARP inspection statistics:
Interface        Packets received  ARP inspection pass   ARP inspection failed
---------------  ----------------  --------------------  ---------------------
ge-0/0/1.0                     7                     5                      2
ge-0/0/2.0                    10                    10                      0
ge-0/0/3.0                    12                    12                      0
```

Meaning The sample output shows the number of ARP packets received and inspected per interface, with a listing of how many packets passed and how many failed the inspection on each interface. The switch compares the ARP requests and replies against the entries in the DHCP snooping database. If a MAC address or IP address in the ARP packet does not match a valid entry in the database, the packet is dropped.

Verifying That MAC Limiting and MAC Move Limiting Are Working Correctly on the Switch

Purpose Verify that MAC limiting and MAC move limiting are working on the switch.

Action Suppose that two packets have been sent from hosts on **ge-0/0/1** and five packets from hosts on **ge-0/0/2**, with both interfaces set to a MAC limit of **4** with the default action **drop**.

Display the MAC addresses learned:

```
user@switch> show ethernet-switching table
```

```
Ethernet-switching table:  7 entries, 6 learned
  VLAN                MAC address         Type         Age    Interfaces

  employee-vlan       *                   Flood          -    ge-0/0/2.0
  employee-vlan       00:05:85:3A:82:77   Learn          0    ge-0/0/1.0
  employee-vlan       00:05:85:3A:82:79   Learn          0    ge-0/0/1.0
  employee-vlan       00:05:85:3A:82:80   Learn          0    ge-0/0/2.0
  employee-vlan       00:05:85:3A:82:81   Learn          0    ge-0/0/2.0
  employee-vlan       00:05:85:3A:82:83   Learn          0    ge-0/0/2.0
  employee-vlan       00:05:85:3A:82:85   Learn          0    ge-0/0/2.0
```

Now suppose packets have been sent from two of the hosts on **ge-0/0/2** after they have been moved to other interfaces more than 5 times in 1 second, with **employee-vlan** set to a MAC move limit of **5** with the default action **drop**.

Display the MAC addresses in the table:

```
user@switch> show ethernet-switching table

Ethernet-switching table:  7 entries, 4 learned
  VLAN                MAC address         Type         Age    Interfaces

  employee-vlan       *                   Flood          -    ge-0/0/2.0
  employee-vlan       00:05:85:3A:82:77   Learn          0    ge-0/0/1.0
  employee-vlan       00:05:85:3A:82:79   Learn          0    ge-0/0/1.0
  employee-vlan       00:05:85:3A:82:80   Learn          0    ge-0/0/2.0
  employee-vlan       00:05:85:3A:82:81   Learn          0    ge-0/0/2.0
  employee-vlan       *                   Flood          -    ge-0/0/2.0
  employee-vlan       *                   Flood          -    ge-0/0/2.0
```

Meaning The first sample output shows that with a MAC limit of **4** for each interface, the fifth MAC address on **ge-0/0/2** was not learned because it exceeded the MAC limit. The second sample output shows that MAC addresses for three of the hosts on **ge-/0/0/2** were not learned, because the hosts had been moved back more than 5 times in one second.

Verifying That Allowed MAC Addresses Are Working Correctly on the Switch

Purpose Verify that allowed MAC addresses are working on the switch.

Action Display the MAC cache information after 5 allowed MAC addresses have been configured on interface **ge-0/0/2**:

```
user@switch> show ethernet-switching table
Ethernet-switching table:  5 entries, 4 learned
  VLAN                MAC address         Type         Age    Interfaces

  employee-vlan       00:05:85:3A:82:80   Learn          0    ge-0/0/2.0
  employee-vlan       00:05:85:3A:82:81   Learn          0    ge-0/0/2.0
  employee-vlan       00:05:85:3A:82:83   Learn          0    ge-0/0/2.0
  employee-vlan       00:05:85:3A:82:85   Learn          0    ge-0/0/2.0
  employee-vlan       *                   Flood          -    ge-0/0/2.0
```

Meaning Because the MAC limit value for this interface has been set to **4**, only 4 of the 5 configured allowed addresses are learned.

Related Documentation

- Example: Configuring DHCP Snooping, DAI , and MAC Limiting on a J-EX Series Switch with Access to a DHCP Server Through a Second Switch on page 2593
- Example: Configuring a DHCP Server Interface as Untrusted to Protect the Switch from Rogue DHCP Server Attacks on page 2579
- Example: Configuring Allowed MAC Addresses to Protect the Switch from DHCP Snooping Database Alteration Attacks on page 2590
- Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks on page 2586
- Example: Configuring MAC Limiting, Including Dynamic and Allowed MAC Addresses, to Protect the Switch from Ethernet Switching Table Overflow Attacks on page 2576
- Example: Configuring MAC Limiting to Protect the Switch from DHCP Starvation Attacks on page 2583
- Configuring Port Security (CLI Procedure) on page 2626
- Configuring Port Security (J-Web Procedure) on page 2627

Example: Configuring MAC Limiting, Including Dynamic and Allowed MAC Addresses, to Protect the Switch from Ethernet Switching Table Overflow Attacks

In an Ethernet switching table overflow attack, an intruder sends so many requests from new MAC addresses that the Ethernet switching table fills up and then overflows, forcing the switch to broadcast all messages.

This example describes how to configure MAC limiting and allowed MAC addresses, two port security features, to protect the switch from Ethernet switching table attacks:

Requirements

This example uses the following hardware and software components:

- One J-EX Series switch
- A DHCP server to provide IP addresses to network devices on the switch

Before you configure specific port security features to mitigate common access-interface attacks, be sure you have:

- Connected the DHCP server to the switch.
- Configured the VLAN **employee-vlan** on the switch. See “Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches” on page 1070.

Overview and Topology

Ethernet LANs are vulnerable to address spoofing and DoS attacks on network devices. This example describes how to protect the switch from an attack on the Ethernet switching table that causes the table to overflow and thus forces the switch to broadcast all messages.

This example shows how to configure port security features on a J-EX4200-24T switch. The switch is connected to a DHCP server.

The setup for this example includes the VLAN **employee-vlan** on the switch. The procedure for creating that VLAN is described in the topic "Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches" on page 1070. That procedure is not repeated here. Figure 68 on page 2577 illustrates the topology for this example.

Figure 68: Network Topology for Basic Port Security

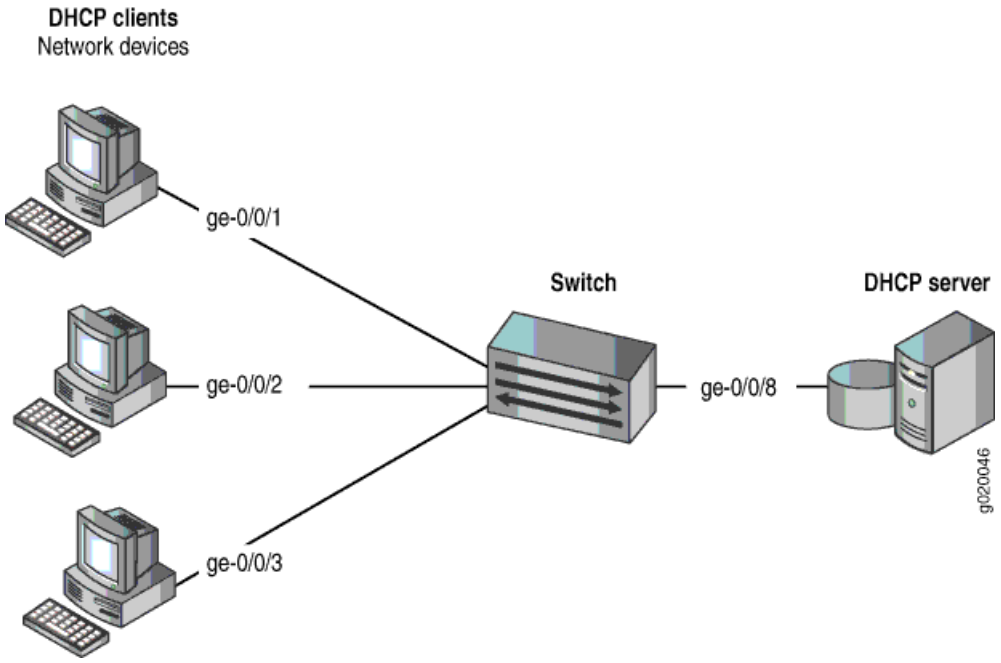

The components of the topology for this example are shown in Table 322 on page 2577.

Table 322: Components of the Port Security Topology

Properties	Settings
Switch hardware	One J-EX4200-24T, 24 ports (8 PoE ports)
VLAN name and ID	**employee-vlan**, tag **20**
VLAN subnets	**192.0.2.16/28** **192.0.2.17** through **192.0.2.30** **192.0.2.31** is subnet's broadcast address
Interfaces in **employee-vlan**	**ge-0/0/1**, **ge-0/0/2**, **ge-0/0/3**, **ge-0/0/8**

Table 322: Components of the Port Security Topology *(continued)*

Properties	Settings
Interface for DHCP server	ge-0/0/8

In this example, use the MAC limit feature to control the total number of MAC addresses that can be added to the Ethernet switching table for the specified interface. Use the allowed MAC addresses feature to ensure that the addresses of network devices whose network access is critical are guaranteed to be included in the Ethernet switching table.

In this example, the switch has already been configured as follows:

- Secure port access is activated on the switch.
- No MAC limit is set on any of the interfaces.
- All access interfaces are untrusted, which is the default setting.

Configuration

To configure MAC limiting and some allowed MAC addresses to protect the switch against Ethernet switching table overflow attacks:

CLI Quick Configuration

To quickly configure MAC limiting and some allowed MAC addresses, copy the following commands and paste them into the switch terminal window:

```
[edit ethernet-switching-options secure-access-port]
set interface ge-0/0/1 mac-limit 4 action drop
set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:80
set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:81
set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:83
set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:85
```

Step-by-Step Procedure

Configure MAC limiting and some allowed MAC addresses:

1. Configure a MAC limit of 4 on **ge-0/0/1** and specify that incoming packets with different addresses be dropped once the limit is exceeded on the interface:

   ```
   [edit ethernet-switching-options secure-access-port]
   user@switch# set interface ge-0/0/1 mac-limit 4 action drop
   ```

2. Configure the allowed MAC addresses on **ge-0/0/2**:

   ```
   [edit ethernet-switching-options secure-access-port]
   user@switch# set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:80
   user@switch# set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:81
   user@switch# set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:83
   user@switch# set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:85
   ```

Results

Check the results of the configuration:

```
[edit ethernet-switching-options secure-access-port]
user@switch# show
interface ge-0/0/1.0 {
   mac-limit 4 action drop;
}
interface ge-0/0/2.0 {
```

```
    allowed-mac [ 00:05:85:3a:82:80 00:05:85:3a:82:81 00:05:85:3a:82:83 00:05:85
    :3a:82:85 ];
  }
```

Verification

To confirm that the configuration is working properly:

- Verifying That MAC Limiting Is Working Correctly on the Switch on page 2579

Verifying That MAC Limiting Is Working Correctly on the Switch

Purpose Verify that MAC limiting is working on the switch.

Action Display the MAC cache information after DHCP requests have been sent from hosts on **ge-0/0/1**, with the interface set to a MAC limit of **4** with the action **drop**, and after four allowed MAC addresses have been configured on interface **ge/0/0/2**:

```
user@switch>  show ethernet-switching table
Ethernet-switching table:  5 entries, 4 learned
  VLAN              MAC address         Type         Age    Interfaces

  employee-vlan     00:05:85:3A:82:71   Learn          0    ge-0/0/1.0
  employee-vlan     00:05:85:3A:82:74   Learn          0    ge-0/0/1.0
  employee-vlan     00:05:85:3A:82:77   Learn          0    ge-0/0/1.0
  employee-vlan     00:05:85:3A:82:79   Learn          0    ge-0/0/1.0
  employee-vlan     *                   Flood          0    ge-0/0/1.0
  employee-vlan     00:05:85:3A:82:80   Learn          0    ge-0/0/2.0
  employee-vlan     00:05:85:3A:82:81   Learn          0    ge-0/0/2.0
  employee-vlan     00:05:85:3A:82:83   Learn          0    ge-0/0/2.0
  employee-vlan     00:05:85:3A:82:85   Learn          0    ge-0/0/2.0
  employee-vlan     *                   Flood          -    ge-0/0/2.0
```

Meaning The sample output shows that with a MAC limit of **4** for the interface, the DHCP request for a fifth MAC address on **ge-0/0/1** was dropped because it exceeded the MAC limit and that only the specified allowed MAC addresses have been learned on the **ge-0/0/2** interface.

Related Documentation

- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Configuring MAC Limiting (CLI Procedure) on page 2635
- Configuring MAC Limiting (J-Web Procedure) on page 2637

Example: Configuring a DHCP Server Interface as Untrusted to Protect the Switch from Rogue DHCP Server Attacks

In a rogue DHCP server attack, an attacker has introduced a rogue server into the network, allowing it to give IP address leases to the network's DHCP clients and to assign itself as the gateway device.

This example describes how to configure a DHCP server interface as untrusted to protect the switch from a rogue DHCP server:

Requirements

This example uses the following hardware and software components:

- One J-EX Series switch
- A DHCP server to provide IP addresses to network devices on the switch

Before you configure an untrusted DHCP server interface to mitigate rogue DHCP server attacks, be sure you have:

- Connected the DHCP server to the switch.
- Enabled DHCP snooping on the VLAN.
- Configured the VLAN **employee-vlan** on the switch. See “Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches” on page 1070.

Overview and Topology

Ethernet LANs are vulnerable to address spoofing and DoS attacks on network devices. This example describes how to protect the switch from rogue DHCP server attacks.

This example shows how to explicitly configure an untrusted interface on a J-EX4200-24T switch. Figure 69 on page 2581 illustrates the topology for this example.

Figure 69: Network Topology for Basic Port Security

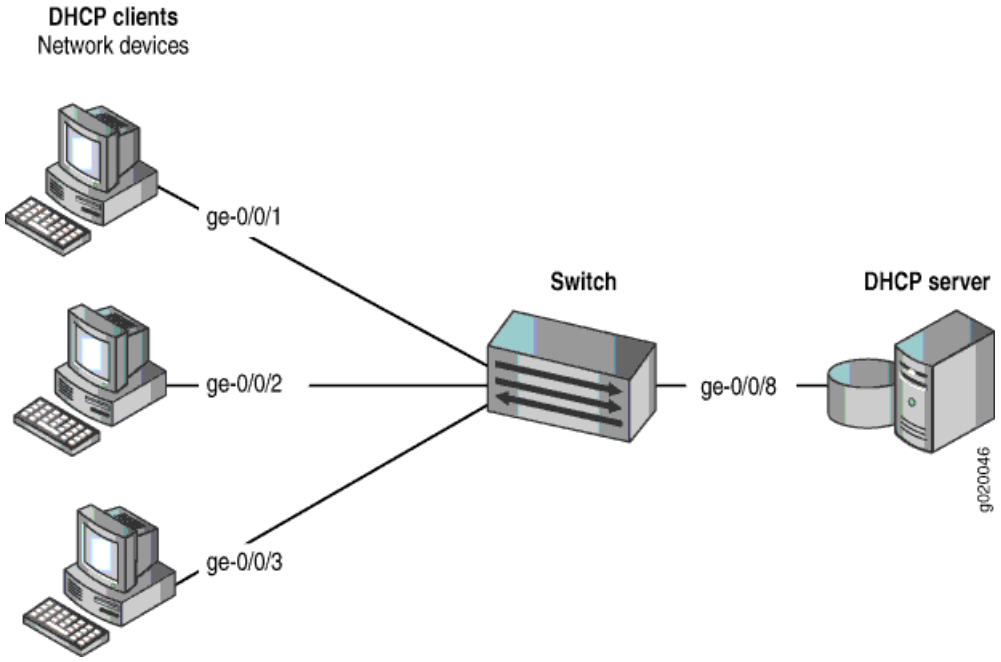

The components of the topology for this example are shown in Table 323 on page 2581.

Table 323: Components of the Port Security Topology

Properties	Settings
Switch hardware	One J-EX4200-24T, 24 ports (8 PoE ports)
VLAN name and ID	**employee-vlan**, tag **20**
VLAN subnets	**192.0.2.16/28** **192.0.2.17** through **192.0.2.30** **192.0.2.31** is the subnet's broadcast address
Interfaces in **employee-vlan**	**ge-0/0/1**, **ge-0/0/2**, **ge-0/0/3**, **ge-0/0/8**
Interface for DHCP server	**ge-0/0/8**

In this example, the switch has already been configured as follows:

- Secure port access is activated on the switch.
- DHCP snooping is enabled on the VLAN **employee-vlan**.
- The interface (port) where the rogue DHCP server has connected to the switch is currently trusted.

Configuration

To configure the DHCP server interface as untrusted because the interface is being used by a rogue DHCP server:

CLI Quick Configuration

To quickly set the rogue DHCP server interface as untrusted, copy the following command and paste it into the switch terminal window:

```
[edit ethernet-switching-options secure-access-port]
set interface ge-0/0/8 no-dhcp-trusted
```

Step-by-Step Procedure

To set the DHCP server interface as untrusted:

Specify the interface (port) from which DHCP responses are not allowed:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface ge-0/0/8 no–dhcp-trusted
```

Results

Check the results of the configuration:

```
[edit ethernet-switching-options secure-access-port]
user@switch# show
interface ge-0/0/8.0 {
   no-dhcp-trusted;
}
```

Verification

To confirm that the configuration is working properly:

- Verifying That the DHCP Server Interface Is Untrusted on page 2582

Verifying That the DHCP Server Interface Is Untrusted

Purpose

Verify that the DHCP server is untrusted.

Action

Send some DHCP requests from network devices (here they are DHCP clients) connected to the switch.

Display the DHCP snooping information when the port on which the DHCP server connects to the switch is not trusted.

```
user@switch> show dhcp snooping binding
```

Meaning

There is no output from the command because no entries are added to the DHCP snooping database.

Related Documentation

- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Enabling a Trusted DHCP Server (CLI Procedure) on page 2632
- Enabling a Trusted DHCP Server (J-Web Procedure) on page 2632

Example: Configuring MAC Limiting to Protect the Switch from DHCP Starvation Attacks

In a DHCP starvation attack, an attacker floods an Ethernet LAN with DHCP requests from spoofed (counterfeit) MAC addresses. The switch's trusted DHCP server or servers cannot keep up with the requests and can no longer assign IP addresses and lease times to legitimate DHCP clients on the switch. Requests from those clients are either dropped or directed to a rogue DHCP server set up by the attacker.

This example describes how to configure MAC limiting, a port security feature, to protect the switch against DHCP starvation attacks:

Requirements

This example uses the following hardware and software components:

- One J-EX Series switch
- A DHCP server to provide IP addresses to network devices on the switch

Before you configure MAC limiting, a port security feature, to mitigate DHCP starvation attacks, be sure you have:

- Connected the DHCP server to the switch.
- Configured the VLAN **employee-vlan** on the switch. See “Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches” on page 1070.

Overview and Topology

Ethernet LANs are vulnerable to address spoofing and DoS attacks on network devices. This example describes how to protect the switch against one common type of attack, a DHCP starvation attack.

This example shows how to configure port security features on a J-EX4200-24T switch that is connected to a DHCP server.

The setup for this example includes the VLAN **employee-vlan** on the switch. The procedure for creating that VLAN is described in the topic “Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches” on page 1070. That procedure is not repeated here. Figure 70 on page 2584 illustrates the topology for this example.

Figure 70: Network Topology for Basic Port Security

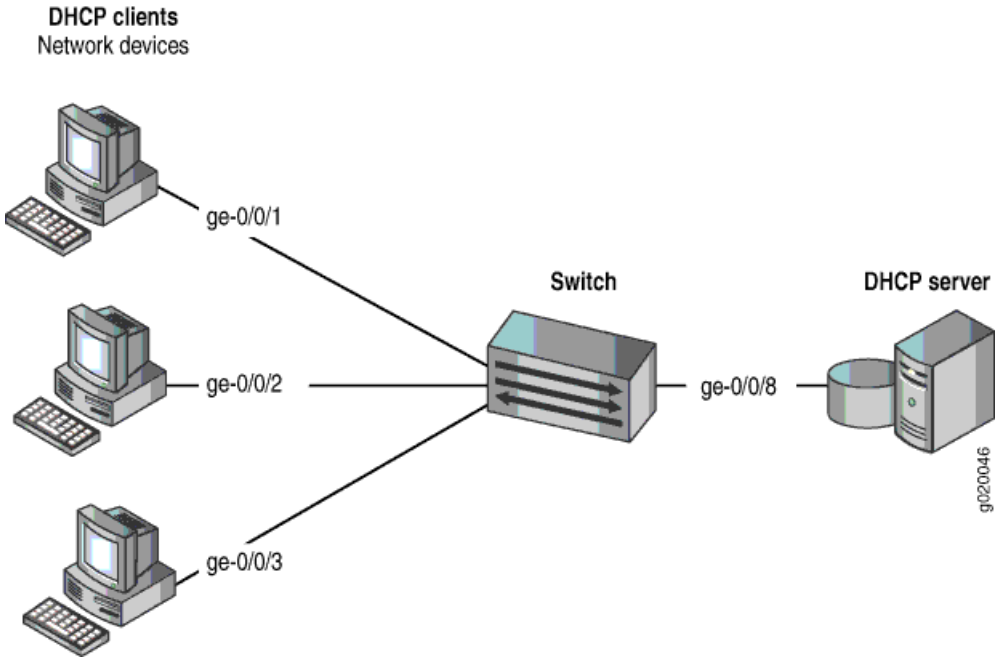

The components of the topology for this example are shown in Table 324 on page 2584.

Table 324: Components of the Port Security Topology

Properties	Settings
Switch hardware	One J-EX4200-24T, 24 ports (8 PoE ports)
VLAN name and ID	**default**
Interfaces in **employee-vlan**	**ge-0/0/1**, **ge-0/0/2**, **ge-0/0/3**, **ge-0/0/8**
Interface for DHCP server	**ge-0/0/8**

In this example, the switch has already been configured as follows:

- Secure port access is activated on the switch.
- No MAC limit is set on any of the interfaces.
- DHCP snooping is disabled on the VLAN **employee-vlan**.
- All access interfaces are untrusted, which is the default setting.

Configuration

To configure the MAC limiting port security feature to protect the switch against DHCP starvation attacks:

CLI Quick Configuration

To quickly configure MAC limiting, copy the following commands and paste them into the switch terminal window:

```
[edit ethernet-switching-options secure-access-port]
set interface ge-0/0/1 mac-limit 3 action drop
set interface ge-0/0/2 mac-limit 3 action drop
```

Step-by-Step Procedure

Configure MAC limiting:

1. Configure a MAC limit of **3** on **ge-0/0/1** and specify that packets with new addresses be dropped if the limit has been exceeded on the interface:

   ```
   [edit ethernet-switching-options secure-access-port]
   user@switch# set interface ge–0/0/1 mac-limit 3 action drop
   ```

2. Configure a MAC limit of **3** on **ge-0/0/2** and specify that packets with new addresses be dropped if the limit has been exceeded on the interface:

   ```
   [edit ethernet-switching-options secure-access-port]
   user@switch# set interface ge-0/0/2 mac-limit 3 action drop
   ```

Results

Check the results of the configuration:

```
[edit ethernet-switching-options secure-access-port]
user@switch# show
interface ge-0/0/1.0 {
  mac-limit 3 action drop;
}
interface ge-0/0/2.0 {
  mac-limit 3 action drop;
}
```

Verification

To confirm that the configuration is working properly:

- Verifying That MAC Limiting Is Working Correctly on the Switch on page 2585

Verifying That MAC Limiting Is Working Correctly on the Switch

Purpose

Verify that MAC limiting is working on the switch.

Action

Send some DHCP requests from network devices (here they are DHCP clients) connected to the switch.

Display the MAC addresses learned when DHCP requests are sent from hosts on **ge-0/0/1** and from hosts on **ge-0/0/2**, with both interfaces set to a MAC limit of **3** with the action **drop**:

```
user@switch> show ethernet-switching table
Ethernet-switching table:  7 entries, 6 learned
  VLAN           MAC address         Type         Age    Interfaces

  default        *                   Flood          -    ge-0/0/2.0
  default        00:05:85:3A:82:77   Learn          0    ge-0/0/1.0
  default        00:05:85:3A:82:79   Learn          0    ge-0/0/1.0
  default        00:05:85:3A:82:80   Learn          0    ge-0/0/1.0
  default        00:05:85:3A:82:81   Learn          0    ge-0/0/2.0
  default        00:05:85:3A:82:83   Learn          0    ge-0/0/2.0
```

```
default        00:05:85:3A:82:85   Learn          0     ge-0/0/2.0
```

Meaning The sample output shows that with a MAC limit of **3** for each interface, the DHCP request for a fourth MAC address on **ge-0/0/2** was dropped because it exceeded the MAC limit.

Because only 3 MAC addresses can be learned on each of the two interfaces, attempted DHCP starvation attacks will fail.

Related Documentation

- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Configuring MAC Limiting (CLI Procedure) on page 2635
- Configuring MAC Limiting (J-Web Procedure) on page 2637

Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks

In an ARP spoofing attack, the attacker associates its own MAC address with the IP address of a network device connected to the switch. Traffic intended for that IP address is now sent to the attacker instead of being sent to the intended destination. The attacker can send faked, or “spoofed,” ARP messages on the LAN.

NOTE: On J-EX Series switches, when dynamic ARP inspection (DAI) is enabled, the switch logs the number of invalid ARP packets that it receives on each interface, along with the sender’s IP and MAC addresses. You can use these log messages to discover ARP spoofing on the network.

This example describes how to configure DHCP snooping and dynamic ARP inspection (DAI), two port security features, to protect the switch against ARP spoofing attacks:

- Requirements on page 2586
- Overview and Topology on page 2587
- Configuration on page 2588
- Verification on page 2589

Requirements

This example uses the following hardware and software components:

- One J-EX Series switch
- A DHCP server to provide IP addresses to network devices on the switch

Before you configure DHCP snooping and DAI, two port security features, to mitigate ARP spoofing attacks, be sure you have:

- Connected the DHCP server to the switch.

- Configured the VLAN **employee-vlan** on the switch.

Overview and Topology

Ethernet LANs are vulnerable to address spoofing and DoS attacks on network devices. This example describes how to protect the switch against one common type of attack, an ARP spoofing attack.

In an ARP spoofing attack, the attacker sends faked ARP messages, thus creating various types of mischief on the LAN—for example, the attacker might launch a man-in-the middle attack.

This example shows how to configure port security features on a J-EX4200-24T switch that is connected to a DHCP server. The setup for this example includes the VLAN **employee-vlan** on the switch. The procedure for creating that VLAN is described in the topic "Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches" on page 1070. That procedure is not repeated here. Figure 71 on page 2587 illustrates the topology for this example.

Figure 71: Network Topology for Basic Port Security

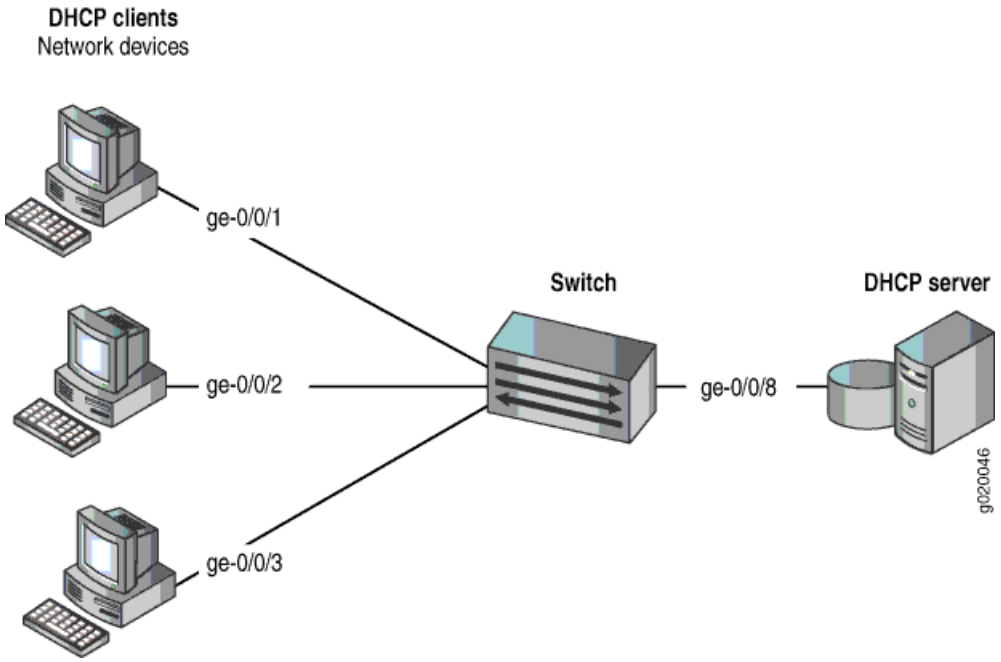

The components of the topology for this example are shown in Table 325 on page 2587.

Table 325: Components of the Port Security Topology

Properties	Settings
Switch hardware	One J-EX4200-24T, 24 ports (8 PoE ports)
VLAN name and ID	**employee-vlan**, tag **20**

Table 325: Components of the Port Security Topology *(continued)*

Properties	Settings
VLAN subnets	**192.0.2.16/28** **192.0.2.17** through **192.0.2.30** **192.0.2.31** is subnet's broadcast address
Interfaces in **employee-vlan**	**ge-0/0/1**,**ge-0/0/2**, **ge-0/0/3**, **ge-0/0/8**
Interface for DHCP server	**ge-0/0/8**

In this example, the switch has already been configured as follows:

- Secure port access is activated on the switch.
- DHCP snooping is disabled on the VLAN **employee-vlan**.
- All access ports are untrusted, which is the default setting.

Configuration

To configure DHCP snooping and dynamic ARP inspection (DAI) to protect the switch against ARP attacks:

CLI Quick Configuration

To quickly configure DHCP snooping and dynamic ARP inspection (DAI), copy the following commands and paste them into the switch terminal window:

```
[edit ethernet-switching-options secure-access-port]
set interface ge-0/0/8 dhcp-trusted
set vlan employee-vlan examine-dhcp
set vlan employee-vlan arp-inspection
```

Step-by-Step Procedure

Configure DHCP snooping and dynamic ARP inspection (DAI) on the VLAN:

1. Set the **ge-0/0/8** interface as trusted:

   ```
   [edit ethernet-switching-options secure-access-port]
   user@switch# set interface ge-0/0/8 dhcp-trusted
   ```

2. Enable DHCP snooping on the VLAN:

   ```
   [edit ethernet-switching-options secure-access-port]
   user@switch# set vlan employee-vlan examine-dhcp
   ```

3. Enable DAI on the VLAN:

   ```
   [edit ethernet-switching-options secure-access-port]
   user@switch# set vlan employee-vlan arp-inspection
   ```

Results

Check the results of the configuration:

```
[edit ethernet-switching-options secure-access-port]
user@switch# show
interface ge-0/0/8.0 {
   dhcp-trusted;
}
vlan employee-vlan {
```

```
    arp-inspection;
    examine-dhcp;
}
```

Verification

To confirm that the configuration is working properly:

- Verifying That DHCP Snooping Is Working Correctly on the Switch on page 2589
- Verifying That DAI Is Working Correctly on the Switch on page 2589

Verifying That DHCP Snooping Is Working Correctly on the Switch

Purpose Verify that DHCP snooping is working on the switch.

Action Send some DHCP requests from network devices (here they are DHCP clients) connected to the switch.

Display the DHCP snooping information when the port on which the DHCP server connects to the switch is trusted. The following output results when requests are sent from the MAC addresses and the server has provided the IP addresses and leases:

```
user@switch> show dhcp snooping binding
DHCP Snooping Information:
MAC Address         IP Address    Lease    Type     VLAN            Interface
-----------------   ----------    -----    ----     ----            ---------
00:05:85:3A:82:77    192.0.2.17    600    dynamic  employee-vlan   ge-0/0/1.0
00:05:85:3A:82:79    192.0.2.18    653    dynamic  employee-vlan   ge-0/0/1.0
00:05:85:3A:82:80    192.0.2.19    720    dynamic  employee-vlan   ge-0/0/2.0
00:05:85:3A:82:81    192.0.2.20    932    dynamic  employee-vlan   ge-0/0/2.0
00:05:85:3A:82:83    192.0.2.21   1230    dynamic  employee-vlan   ge-0/0/2.0
00:05:85:27:32:88    192.0.2.22   3200    dynamic  employee-vlan   ge-0/0/3.0
```

Meaning When the interface on which the DHCP server connects to the switch has been set to trusted, the output (see preceding sample) shows, for each MAC address, the assigned IP address and lease time—that is, the time, in seconds, remaining before the lease expires.

Verifying That DAI Is Working Correctly on the Switch

Purpose Verify that DAI is working on the switch.

Action Send some ARP requests from network devices connected to the switch.

Display the DAI information:

```
user@switch> show arp inspection statistics
ARP inspection statistics:
Interface        Packets received  ARP inspection pass   ARP inspection failed
--------------   ---------------   -------------------- ---------------------
ge-0/0/1.0                     7                     5                     2
ge-0/0/2.0                    10                    10                     0
ge-0/0/3.0                    12                    12                     0
```

Meaning The sample output shows the number of ARP packets received and inspected per interface, with a listing of how many packets passed and how many failed the inspection on each interface. The switch compares the ARP requests and replies against the entries in the DHCP snooping database. If a MAC address or IP address in the ARP packet does not match a valid entry in the database, the packet is dropped.

Related Documentation

- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Enabling DHCP Snooping (CLI Procedure) on page 2630
- Enabling DHCP Snooping (J-Web Procedure) on page 2631
- Enabling Dynamic ARP Inspection (CLI Procedure) on page 2633
- Enabling Dynamic ARP Inspection (J-Web Procedure) on page 2634

Example: Configuring Allowed MAC Addresses to Protect the Switch from DHCP Snooping Database Alteration Attacks

In one type of attack on the DHCP snooping database, an intruder introduces a DHCP client on an untrusted access interface with a MAC address identical to that of a client on another untrusted interface. The intruder then acquires the DHCP lease of that other client, thus changing the entries in the DHCP snooping table. Subsequently, what would have been valid ARP requests from the legitimate client are blocked.

This example describes how to configure allowed MAC addresses, a port security feature, to protect the switch from DHCP snooping database alteration attacks:

Requirements

This example uses the following hardware and software components:

- One J-EX Series switch
- A DHCP server to provide IP addresses to network devices on the switch

Before you configure specific port security features to mitigate common access-inteface attacks, be sure you have:

- Connected the DHCP server to the switch.
- Configured the VLAN **employee-vlan** on the switch. See “Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches” on page 1070.

Overview and Topology

Ethernet LANs are vulnerable to address spoofing and DoS attacks on network devices. This example describes how to protect the switch from an attack on the DHCP snooping database that alters the MAC addresses assigned to some clients.

This example shows how to configure port security features on a J-EX4200-24T switch that is connected to a DHCP server.

The setup for this example includes the VLAN **employee-vlan** on the switch. The procedure for creating that VLAN is described in the topic "Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches" on page 1070. That procedure is not repeated here. Figure 72 on page 2591 illustrates the topology for this example.

Figure 72: Network Topology for Basic Port Security

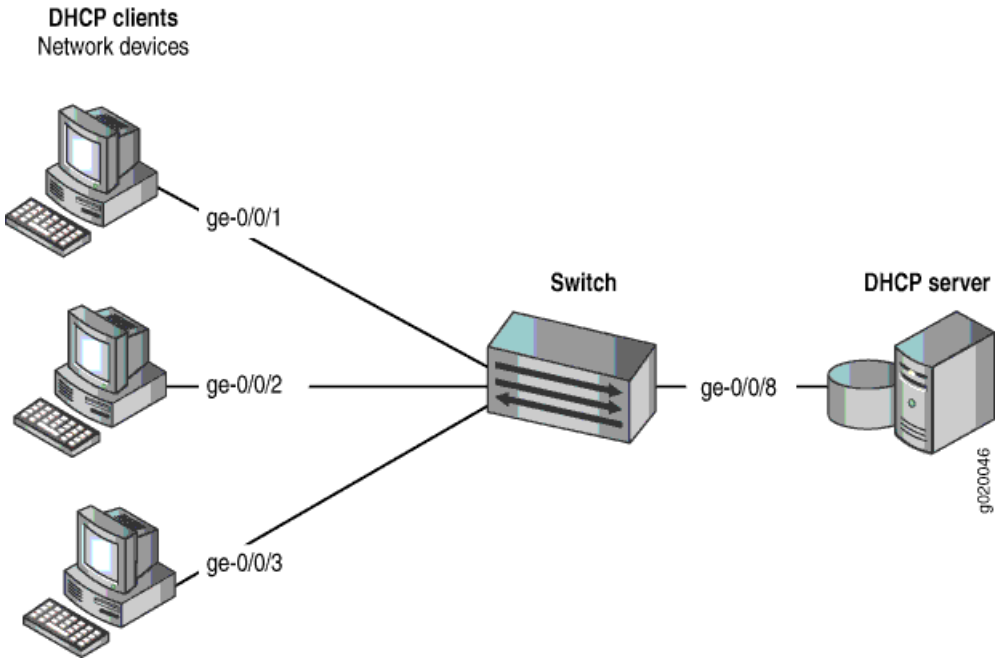

The components of the topology for this example are shown in Table 326 on page 2591.

Table 326: Components of the Port Security Topology

Properties	Settings
Switch hardware	One J-EX4200-24T, 24 ports (8 PoE ports)
VLAN name and ID	**employee-vlan**, tag **20**
VLAN subnets	**192.0.2.16/28** **192.0.2.17** through **192.0.2.30** **192.0.2.31** is subnet's broadcast address
Interfaces in **employee-vlan**	**ge-0/0/1**, **ge-0/0/2**, **ge-0/0/3**, **ge-0/0/8**

Table 326: Components of the Port Security Topology *(continued)*

Properties	Settings
Interface for DHCP server	**ge-0/0/8**

In this example, the switch has already been configured as follows:

- Secure port access is activated on the switch.
- DHCP snooping is enabled on the VLAN **employee-vlan**.
- All access ports are untrusted, which is the default setting.

Configuration

To configure allowed MAC addresses to protect the switch against DHCP snooping database alteration attacks:

CLI Quick Configuration

To quickly configure some allowed MAC addresses on an interface, copy the following commands and paste them into the switch terminal window:

```
[edit ethernet-switching-options secure-access-port]
set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:80
set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:81
set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:83
set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:85
set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:88
```

Step-by-Step Procedure

To configure some allowed MAC addresses on an interface:

Configure the five allowed MAC addresses on an interface:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:80
user@switch# set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:81
user@switch# set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:83
user@switch# set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:85
user@switch# set interface ge-0/0/2 allowed-mac 00:05:85:3A:82:88
```

Results

Check the results of the configuration:

```
[edit ethernet-switching-options secure-access-port]
user@switch# show
interface ge-0/0/2.0 {
   allowed-mac [ 00:05:85:3a:82:80 00:05:85:3a:82:81 00:05:85:3a:82:83 00:05:85
   :3a:82:85 00:05:85:3a:82:88 ];
}
```

Verification

To confirm that the configuration is working properly:

- Verifying That Allowed MAC Addresses Are Working Correctly on the Switch on page 2592

Verifying That Allowed MAC Addresses Are Working Correctly on the Switch

Purpose

Verify that allowed MAC addresses are working on the switch.

Action Display the MAC cache information:

```
user@switch> show ethernet-switching table
Ethernet-switching table:  6 entries, 5 learned
  VLAN                MAC address         Type         Age    Interfaces

  employee-vlan       00:05:85:3A:82:80   Learn          0    ge-0/0/2.0
  employee-vlan       00:05:85:3A:82:81   Learn          0    ge-0/0/2.0
  employee-vlan       00:05:85:3A:82:83   Learn          0    ge-0/0/2.0
  employee-vlan       00:05:85:3A:82:85   Learn          0    ge-0/0/2.0
  employee-vlan       00:05:85:3A:82:88   Learn          0    ge-0/0/2.0
  employee-vlan       *                   Flood          -    ge-0/0/2.0
```

Meaning The output shows that the five MAC addresses configured as allowed MAC addresses have been learned and are displayed in the MAC cache. The last MAC address in the list, one that had not been configured as allowed, has not been added to the list of learned addresses.

Related Documentation

- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Configuring MAC Limiting (CLI Procedure) on page 2635
- Configuring MAC Limiting (J-Web Procedure) on page 2637

Example: Configuring DHCP Snooping, DAI , and MAC Limiting on a J-EX Series Switch with Access to a DHCP Server Through a Second Switch

You can configure DHCP snooping, dynamic ARP inspection (DAI), and MAC limiting on the access interfaces of J-EX Series switches to protect the switch and the Ethernet LAN against address spoofing and Layer 2 denial-of-service (DoS) attacks. To obtain those basic settings, you can use the switch's default configuration for port security, configure the MAC limit, and enable DHCP snooping and DAI on a VLAN. You can configure those features when the DHCP server is connected to a different switch from the one to which the DHCP clients (network devices) are connected.

This example describes how to configure port security features on a J-EX Series switch whose hosts obtain IP addresses and lease times from a DHCP server connected to a second switch:

- Requirements on page 2594
- Overview and Topology on page 2594
- Configuring a VLAN, Interfaces, and Port Security Features on Switch 1 on page 2596
- Configuring a VLAN and Interfaces on Switch 2 on page 2598
- Verification on page 2599

Requirements

This example uses the following hardware and software components:

- One J-EX4200-24T switch—“Switch 1” in this example.
- An additional J-EX Series switch—”Switch 2” in this example. You will not configure port security on this switch.
- A DHCP server connected to Switch 2. You will use the server to provide IP addresses to network devices connected to Switch 1.
- At least two network devices (hosts) that you will connect to access interfaces on Switch 1. These devices will be DHCP clients.

Before you configure DHCP snooping, DAI, and MAC limiting port security features, be sure you have:

- Connected the DHCP server to Switch 2.
- Configured the VLAN **employee-vlan** on the switch. See “Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches” on page 1070.

Overview and Topology

Ethernet LANs are vulnerable to address spoofing and DoS attacks on network devices. To protect the devices from such attacks, you can configure:

- DHCP snooping to validate DHCP server messages
- DAI to protect against ARP spoofing
- MAC limiting to constrain the number of MAC addresses the switch adds to its MAC address cache

This example shows how to configure these port security features on a J-EX4200 switch, which is Switch 1 in this example. Switch 1 is connected to a switch that is not configured with port security features. That second switch (Switch 2) is connected to a DHCP server. (See Figure 73 on page 2595.) Network devices (hosts) that are connected to Switch 1 will send requests for IP addresses (that is, the devices will be DHCP clients). Those requests will be transmitted from Switch 1 to Switch 2 and then to the DHCP server connected to Switch 2. Responses to the requests will be transmitted along the reverse path of the one followed by the requests.

The setup for this example includes the VLAN **employee-vlan** on both switches.

Figure 73 on page 2595 shows the network topology for the example.

Figure 73: Network Topology for Port Security Setup with Two Switches on the Same VLAN

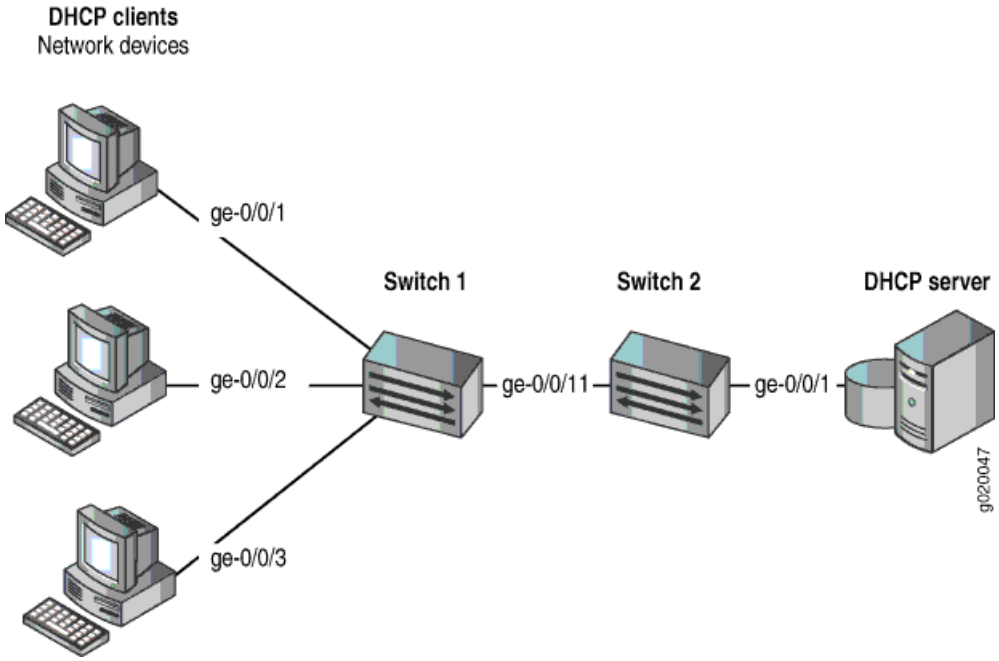

The components of the topology for this example are shown in Table 327 on page 2595.

Table 327: Components of Port Security Setup on Switch 1 with a DHCP Server Connected to Switch 2

Properties	Settings
Switch hardware	One J-EX4200-24T (Switch 1), and an additional J-EX Series switch (Switch 2)
VLAN name and ID	**employee-vlan**, tag **20**
VLAN subnets	**192.0.2.16/28** **192.0.2.17** through **192.0.2.30** **192.0.2.31** is subnet's broadcast address
Trunk interface on both switches	**ge-0/0/11**
Access interfaces on Switch 1	**ge-0/0/1**, **ge-0/0/2**, and **ge-0/0/3**
Access interface on Switch 2	**ge-0/0/1**
Interface for DHCP server	**ge-0/0/1** on Switch 2

Switch 1 is initially configured with the default port security setup. In the default configuration on the switch:

- Secure port access is activated on the switch.
- The switch does not drop any packets, which is the default setting.
- DHCP snooping and dynamic ARP inspection (DAI) are disabled on all VLANs.
- All access interfaces are untrusted and trunk interfaces are trusted; these are the default settings.

In the configuration tasks for this example, you configure a VLAN on both switches.

In addition to configuring the VLAN, you enable DHCP snooping on Switch 1. In this example, you will also enable DAI and a MAC limit of **5** on Switch 1.

Because the interface that connects Switch 2 to Switch 1 is a trunk interface, you do not have to configure this interface to be trusted. As noted above, trunk interfaces are automatically trusted, so DHCP messages coming from the DHCP server to Switch 2 and then on to Switch 1 are trusted.

Configuring a VLAN, Interfaces, and Port Security Features on Switch 1

To configure a VLAN, interfaces, and port security features on Switch 1:

CLI Quick Configuration

To quickly configure a VLAN, interfaces, and port security features, copy the following commands and paste them into the switch terminal window:

```
[edit]
set ethernet-switching-options secure-access-port interface ge-0/0/1 mac–limit 5
set ethernet-switching-options secure-access-port vlan employee-vlan arp–inspection
set ethernet-switching-options secure-access-port vlan employee-vlan examine–dhcp
set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan members 20
set interfaces ge-0/0/2 unit 0 family ethernet-switching vlan members 20
set interfaces ge-0/0/3 unit 0 family ethernet-switching vlan members 20
set interfaces ge-0/0/11 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/11 unit 0 family ethernet-switching vlan members 20
set vlans employee–vlan vlan-id 20
```

Step-by-Step Procedure

To configure MAC limiting, a VLAN, and interfaces on Switch 1 and enable DAI and DHCP on the VLAN:

1. Configure the VLAN **employee-vlan** with VLAN ID **20**:

   ```
   [edit vlans]
   user@switch1# set employee-vlan vlan-id 20
   ```

2. Configure an interface on Switch 1 as a trunk interface:

   ```
   [edit interfaces]
   user@switch1# set ge-0/0/11 unit 0 family ethernet-switching port-mode trunk
   ```

3. Associate the VLAN with interfaces **ge-0/0/1**, **ge-0/0/2**, **ge-0/0/3**, and **ge-0/0/11**:

   ```
   [edit interfaces]
   user@switch1# set ge-0/0/1 unit 0 family ethernet-switching vlan members 20
   user@switch1# set ge-0/0/2 unit 0 family ethernet-switching vlan members 20
   user@switch1# set ge-0/0/3 unit 0 family ethernet-switching vlan members 20
   user@switch1# set ge-0/0/11 unit 0 family ethernet-switching vlan members 20
   ```

4. Enable DHCP snooping on the VLAN:

   ```
   [edit ethernet-switching-options secure-access-port]
   ```

```
user@switch1# set vlan employee-vlan examine-dhcp
```

5. Enable DAI on the VLAN:

```
[edit ethernet-switching-options secure-access-port]
user@switch1# set vlan employee-vlan arp-inspection
```

6. Configure a MAC limit of **5** on **ge-0/0/1** and use the default action, drop (packets with new addresses are dropped if the limit has been exceeded):

```
[edit ethernet-switching-options secure-access-port]
user@switch1# set interface ge-0/0/1 mac-limit 5
```

Results Display the results of the configuration:

```
[edit]
user@switch1# show
ethernet-switching-options {
  secure-access-port {
    interface ge-0/0/1.0{
      mac-limit 5 action drop;
    }
    vlan employee-vlan {
      arp-inspection;
      examine-dhcp;
    }
  }
}
interfaces {
  ge-0/0/1 {
    unit 0 {
      family ethernet-switching {
        vlan {
          members 20;
        }
      }
    }
  }
  ge-0/0/2 {
    unit 0 {
      family ethernet-switching {
        vlan {
          members 20;
        }
      }
    }
  }
  ge-0/0/3 {
    unit 0 {
      family ethernet-switching {
        vlan {
          members 20;
        }
      }
    }
  }
```

```
    ge-0/0/11 {
      unit 0 {
        family ethernet-switching {
          port-mode trunk;
          vlan {
            members 20;
          }
        }
      }
    }
}
vlans {
    employee-vlan {
      vlan-id 20;
    }
}
```

Configuring a VLAN and Interfaces on Switch 2

To configure the VLAN and interfaces on Switch 2:

CLI Quick Configuration

To quickly configure the VLAN and interfaces on Switch 2, copy the following commands and paste them into the switch terminal window:

```
[edit]
set interfaces ge-0/0/1 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/11 unit 0 family ethernet-switching port-mode trunk
set vlans employee-vlan vlan-id 20
```

Step-by-Step Procedure

To configure the VLAN and interfaces on Switch 2:

1. Configure an interface on Switch 2 as a trunk interface:

```
[edit interfaces]
user@switch2# set ge-0/0/11 unit 0 ethernet-switching port-mode trunk
```

2. Associate the VLAN with interfaces **ge-0/0/1** and **ge-0/0/11**:

```
[edit interfaces]
user@switch2# set ge-0/0/1 unit 0 family ethernet-switching vlan members 20
user@switch2# set ge-0/0/11 unit 0 family ethernet-switching vlan members 20
```

Results

Display the results of the configuration:

```
[edit]
user@switch2# show
interfaces {
    ge-0/0/1 {
      unit 0 {
        family ethernet-switching {
          vlan {
            members 20;
          }
        }
      }
    }
    ge-0/0/11 {
```

```
        unit 0 {
          family ethernet-switching {
            port-mode trunk;
            vlan {
              members 20;
            }
          }
        }
      }
    }
    vlans {
      employee-vlan {
        vlan-id 20;
      }
    }
```

Verification

To confirm that the configuration is working properly:

- Verifying That DHCP Snooping Is Working Correctly on Switch 1 on page 2599
- Verifying That DAI Is Working Correctly on Switch 1 on page 2599
- Verifying That MAC Limiting Is Working Correctly on Switch 1 on page 2600

Verifying That DHCP Snooping Is Working Correctly on Switch 1

Purpose Verify that DHCP snooping is working on Switch 1.

Action Send some DHCP requests from network devices (here they are DHCP clients) connected to the switch.

Display the DHCP snooping information when the interface through which Switch 2 sends the DHCP server replies to clients connected to Switch 1 is trusted. The server has provided the IP addresses and leases:

```
user@switch1> show dhcp snooping binding
DHCP Snooping Information:
MAC Address         IP Address    Lease    Type     VLAN          Interface
-----------------   ----------    -----    ----     ----          ---------
00:05:85:3A:82:77    192.0.2.17    600    dynamic  employee-vlan   ge-0/0/1.0
00:05:85:3A:82:79    192.0.2.18    653    dynamic  employee-vlan   ge-0/0/1.0
00:05:85:3A:82:80    192.0.2.19    720    dynamic  employee-vlan   ge-0/0/1.0
00:05:85:3A:82:81    192.0.2.20    932    dynamic  employee-vlan   ge-0/0/1.0
00:05:85:3A:82:83    192.0.2.21   1230    dynamic  employee-vlan   ge-0/0/1.0
00:05:85:3A:82:90    192.0.2.20    932    dynamic  employee-vlan   ge-0/0/2.0
00:05:85:3A:82:91    192.0.2.21   1230    dynamic  employee-vlan   ge-0/0/3.0
```

Meaning The output shows, for each MAC address, the assigned IP address and lease time—that is, the time, in seconds, remaining before the lease expires.

Verifying That DAI Is Working Correctly on Switch 1

Purpose Verify that DAI is working on Switch 1.

Action Send some ARP requests from network devices connected to the switch.

Display the DAI information:

```
user@switch1> show arp inspection statistics
ARP inspection statistics:
Interface    Packets received  ARP inspection pass  ARP inspection failed
----------   ----------------  -------------------  ---------------------
ge-0/0/1.0                 7                    5                      2
ge-0/0/2.0                10                   10                      0
ge-0/0/3.0                18                   15                      3
```

Meaning The sample output shows the number of ARP packets received and inspected per interface, with a listing of how many packets passed and how many failed the inspection on each interface. The switch compares the ARP requests and replies against the entries in the DHCP snooping database. If a MAC address or IP address in the ARP packet does not match a valid entry in the database, the packet is dropped.

Verifying That MAC Limiting Is Working Correctly on Switch 1

Purpose Verify that MAC limiting is working on Switch 1.

Action Display the MAC addresses that are learned when DHCP requests are sent from hosts on **ge-0/0/1**:

```
user@switch1> show ethernet-switching table

Ethernet-switching table:  6 entries, 5 learned
  VLAN                MAC address         Type         Age    Interfaces

  employee-vlan       00:05:85:3A:82:77   Learn          0    ge-0/0/1.0
  employee-vlan       00:05:85:3A:82:79   Learn          0    ge-0/0/1.0
  employee-vlan       00:05:85:3A:82:80   Learn          0    ge-0/0/1.0
  employee-vlan       00:05:85:3A:82:81   Learn          0    ge-0/0/1.0
  employee-vlan       00:05:85:3A:82:83   Learn          0    ge-0/0/1.0
  employee-vlan       *                   Flood          -    ge-0/0/1.0
```

Meaning The sample output shows that five MAC addresses have been learned for interface **ge-0/0/1**, which corresponds to the MAC limit of **5** set in the configuration. The last line of the output shows that a sixth MAC address request was dropped, as indicated by the asterisk (*) in the **MAC address** column.

Related Documentation

- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Configuring Port Security (CLI Procedure) on page 2626
- Configuring Port Security (J-Web Procedure) on page 2627

Example: Configuring IP Source Guard with Other J-EX Series Switch Features to Mitigate Address-Spoofing Attacks on Untrusted Access Interfaces

Ethernet LAN switches are vulnerable to attacks that involve spoofing (forging) of source IP addresses or source MAC addresses. These spoofed packets are sent from hosts

connected to untrusted access interfaces on the switch. You can enable the IP source guard port security feature on J-EX Series switches to mitigate the effects of such attacks. If IP source guard determines that a source IP address and a source MAC address in a binding in an incoming packet are not valid, the switch does not forward the packet.

You can use IP source guard in combination with other J-EX Series switch features to mitigate address-spoofing attacks on untrusted access interfaces. This example shows two configuration scenarios:

Requirements

This example uses the following hardware and software components:

- A J-EX4200-24T switch
- A DHCP server to provide IP addresses to network devices on the switch
- A RADIUS server to provide 802.1X authentication

Before you configure IP source guard for these scenarios, be sure you have:

- Connected the DHCP server to the switch.
- Connected the RADIUS server and configured user authentication on the RADIUS server. See “Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch” on page 2267.
- Configured the VLANs on the switch. See “Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches” on page 1070 for detailed information about configuring VLANs.

Overview and Topology

IP source guard checks the IP source address and MAC source address in a packet sent from a host attached to an untrusted access interface on the switch. If IP source guard determines that the packet header contains an invalid source IP address or source MAC address, it ensures that the switch does not forward the packet—that is, the packet is discarded.

When you configure IP source guard, you enable on it on one or more VLANs. IP source guard applies its checking rules to untrusted access interfaces on those VLANs. By default, on J-EX Series switches, access interfaces are untrusted and trunk interfaces are trusted. IP source guard does not check packets that have been sent to the switch by devices connected to either trunk interfaces or trusted access interfaces—that is, interfaces

configured with **dhcp-trusted** so that a DHCP server can be connected to that interface to provide dynamic IP addresses.

IP source guard obtains information about IP-address/MAC-address/VLAN bindings from the DHCP snooping database. It causes the switch to validate incoming IP packets against the entries in that database.

The topology for this example includes a J-EX4200-24P switch, a connection to a DHCP server, and a connection to a RADIUS server for user authentication.

NOTE: The 802.1X user authentication applied in this example is for single supplicants. Single-secure supplicant mode and multiple supplicant mode do not work with IP source guard. For more information about 802.1X authentication, see "Understanding Authentication on J-EX Series Switches" on page 2248.

In the first example configuration, two clients (network devices) are connected to an access switch. You configure IP source guard and 802.1X user authentication, in combination with two access port security features: DHCP snooping and dynamic ARP inspection (DAI). This setup is designed to protect the switch from IP attacks such as "ping of death" attacks, DHCP starvation, and ARP spoofing.

In the second example configuration, the switch is configured for 802.1X user authentication. If the client fails authentication, the switch redirects the client to a guest VLAN that allows this client to access a set of restricted network features. You configure IP source guard on the guest VLAN to mitigate effects of source IP spoofing.

NOTE: Control-plane rate limiting is achieved by restricting CPU control-plane protection. It can be used in conjunction with storm control (see "Understanding Storm Control on J-EX Series Switches" on page 2511) to limit data-plane activity.

TIP: You can set the ip-source-guard flag in the traceoptions statement for debugging purposes.

Configuring IP Source Guard with 802.1X Authentication, DHCP Snooping, and Dynamic ARP Inspection

CLI Quick Configuration

To quickly configure IP source guard with 802.1X authentication and with other access port security features, copy the following commands and paste them into the switch terminal window:

```
[edit]
set ethernet-switching-options secure-access-port interface ge-0/0/24 dhcp-trusted
set ethernet-switching-options secure-access-port vlan data examine-dhcp
set ethernet-switching-options secure-access-port vlan data arp-inspection
set ethernet-switching-options secure-access-port vlan data ip-source-guard
```

```
set interfaces ge-0/0/0 unit 0 family ethernet-switching vlan members data
set interfaces ge-0/0/1 unit 0 family ethernet-switching vlan members data
set interfaces ge-0/0/24 unit 0 family ethernet-switching vlan members data
set protocols lldp-med interface ge-0/0/0.0
set protocols dot1x authenticator authentication-profile-name profile52
set protocols dot1x authenticator interface ge-0/0/0.0 supplicant single
set protocols lldp-med interface ge-0/0/1.0
set protocols dot1x authenticator interface ge-0/0/1.0 supplicant single
```

Step-by-Step Procedure

To configure IP source guard with 802.1X authentication and various port security features:

1. Configure the interface on which the DHCP server is connected to the switch as a trusted interface and add that interface to the data VLAN:

   ```
   [edit ethernet-switching-options]
   user@switch# set secure-access-port interface ge-0/0/24 dhcp-trusted
   user@switch# set set ge-0/0/24 unit 0 family ethernet-switching vlan members data
   ```

2. Associate two interfaces with the data VLAN:

   ```
   [edit interfaces]
   user@switch# set ge-0/0/0 unit 0 family ethernet-switching vlan members data
   user@switch# set ge-0/0/1 unit 0 family ethernet-switching vlan members data
   ```

3. Configure 802.1X user authentication and LLDP-MED on the two interfaces that you associated with the data VLAN:

   ```
   [edit protocols]
   user@switch# set lldp-med interface ge-0/0/0.0
   user@switch# set dot1x authenticator authentication-profile-name profile52
   user@switch# set dot1x authenticator interface ge-0/0/0.0 supplicant single
   user@switch# set lldp-med interface ge-0/0/1.0
   user@switch# set dot1x authenticator interface ge-0/0/1.0 supplicant single
   ```

4. Configure three access port security features—DHCP snooping, dynamic ARP inspection (DAI), and IP source guard—on the data VLAN:

   ```
   [edit ethernet-switching-options]
   user@switch# set secure-access-port vlan data examine-dhcp
   user@switch# set secure-access-port vlan data arp-inspection
   user@switch# set secure-access-port vlan data ip-source-guard
   ```

Results

Check the results of the configuration:

```
[edit ethernet-switching-options]
secure-access-port {
  interface ge-0/0/24.0 {
    dhcp-trusted;
  }
  vlan data {
    arp-inspection;
    examine-dhcp;
    ip-source-guard;
  }
}

[edit interfaces]
ge-0/0/0 {
  unit 0 {
```

```
      family ethernet-switching {
        vlan {
          members data;
        }
      }
    }
  }
  ge-0/0/1 {
    unit 0 {
      family ethernet-switching {
        vlan {
          members data;
        }
      }
    }
  }
  ge-0/0/24 {
    unit 0 {
      family ethernet-switching {
        vlan {
          members data;
        }
      }
    }
  }

  [edit protocols]
  lldp-med {
    interface ge-0/0/14.0;
    interface ge-0/0/0.0;
    interface ge-0/0/1.0;
  }
  dot1x {
    authenticator {
      authentication-profile-name profile52;
    }
    interface {
      ge-0/0/0.0 {
        supplicant single;
      }
      ge-0/0/1.0 {
        supplicant single;
      }
      ge-0/0/14.0 {
        supplicant single;
      }
    }
  }
```

Configuring IP Source Guard on a Guest VLAN

CLI Quick Configuration

To quickly configure IP source guard on a guest VLAN, copy the following commands and paste them into the switch terminal window:

```
[edit]
set ethernet-switching-options secure-access-port interface ge-0/0/24 dhcp-trusted
set interfaces ge-0/0/24 unit 0 family ethernet-switching vlan members employee
```

```
set ethernet-switching-options secure-access-port vlan employee examine-dhcp
set ethernet-switching-options secure-access-port vlan employee ip-source-guard
set ethernet-switching-options secure-access-port interface ge-0/0/0 static-ip 11.1.1.1 mac
00:11:11:11:11:11 vlan employee
set ethernet-switching-options secure-access-port interface ge-0/0/1 static-ip 11.1.1.2 mac
00:22:22:22:22:22 vlan employee
set interfaces ge-0/0/0 unit 0 family ethernet-switching port-mode access
set interfaces ge-0/0/1 unit 0 family ethernet-switching port-mode access
set protocols dot1x authenticator authentication-profile-name profile52
set protocols dot1x authenticator interface ge-0/0/0 supplicant single
set protocols dot1x authenticator interface ge-0/0/0 guest-vlan employee
set protocols dot1x authenticator interface ge-0/0/0 supplicant-timeout 2
set protocols dot1x authenticator interface ge-0/0/1 supplicant single
set protocols dot1x authenticator interface ge-0/0/1 guest-vlan employee
set protocols dot1x authenticator interface ge-0/0/1 supplicant-timeout 2
set vlans employee vlan-id 300
```

Step-by-Step Procedure

To configure IP source guard on a guest VLAN:

1. Configure the interface on which the DHCP server is connected to the switch as a trusted interface and add that interface to the **employee** VLAN:

   ```
   [edit ethernet-switching-options]
   user@switch# set secure-access-port interface ge-0/0/24 dhcp-trusted
   user@switch# set ge-0/0/24 unit 0 family ethernet-switching vlan members employee
   ```

2. Configure two interfaces for the access port mode:

   ```
   [edit interfaces]
   user@switch# set ge-0/0/0 unit 0 family ethernet-switching port-mode access
   user@switch# set ge-0/0/1 unit 0 family ethernet-switching port-mode access
   ```

3. Configure DHCP snooping and IP source guard on the **employee** VLAN:

   ```
   [edit ethernet-switching-options]
   user@switch# set secure-access-port vlan employee examine-dhcp
   user@switch# set secure-access-port vlan employee ip-source-guard
   ```

4. Configure a static IP address on each of two interfaces on the **employee** VLAN (optional):

   ```
   [edit ethernet-switching-options]
   user@switch# set secure-access-port interface ge-0/0/0 static-ip 11.1.1.1 mac
   00:11:11:11:11:11 vlan employee


   [edit ethernet-switching-options]
   user@switch# set secure-access-port interface ge-0/0/1 static-ip 11.1.1.2 mac
   00:22:22:22:22:22 vlan employee
   ```

5. Configure 802.1X user authentication:

   ```
   [edit protocols]
   user@switch# set dot1x authenticator authentication-profile-name profile52
   user@switch# set dot1x authenticator interface ge-0/0/0 supplicant single
   user@switch# set dot1x authenticator interface ge-0/0/1 supplicant single
   user@switch# set dot1x authenticator interface ge-0/0/0 supplicant-timeout 2
   user@switch# set dot1x authenticator interface ge-0/0/1 supplicant-timeout 2
   ```

6. Set the VLAN ID for the **employee** VLAN:

   ```
   [edit vlans]
   user@switch# set employee vlan-id 100
   ```

Results Check the results of the configuration:

```
[edit protocols]
  dot1x {
    authenticator {
      authentication-profile-name profile52;
    }
    interface {
      ge-0/0/0.0 {
        guest-vlan employee;
        supplicant single;
        supplicant-timeout 2;
      }
      ge-0/0/1.0 {
        guest-vlan employee;
        supplicant single;
        supplicant-timeout 2;
      }
    }
  }
}

[edit vlans]
employee {
  vlan-id 100;
}

[edit interfaces]
ge-0/0/0 {
  unit 0 {
    family ethernet-switching {
      port-mode access;
    }
  }
}
ge-0/0/1 {
  unit 0 {
    family ethernet-switching {
      port-mode access;
    }
  }
}
ge-0/0/24 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members employee;
      }
    }
  }
}

[edit ethernet-switching-options]
secure-access-port {
  interface ge-0/0/0.0 {
    static-ip 11.1.1.1 vlan employee mac 00:11:11:11:11:11;
```

```
    }
    interface ge-0/0/1.0 {
      static-ip 11.1.1.2 vlan employee mac 00:22:22:22:22:22;
    }
    interface ge-0/0/24.0 {
      dhcp-trusted;
    }
    vlan employee {
      examine-dhcp;
      ip-source-guard;
    }
  }
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying That 802.1X User Authentication Is Working on the Interface on page 2607
- Verifying the VLAN Association with the Interface on page 2607
- Verifying That DHCP Snooping and IP Source Guard Are Working on the VLAN on page 2607

Verifying That 802.1X User Authentication Is Working on the Interface

Purpose Verify that the 802.1X configuration is working on the interface.

Action Use the **show dot1x interface** command to view the 802.1X details.

Meaning The **Supplicant mode** output field displays the configured administrative mode for each interface.

Verifying the VLAN Association with the Interface

Purpose Verity interface states and VLAN memberships.

Action Use the **show ethernet-switching interfaces** command to view the Ethernet switching table entries.

Meaning The field **VLAN members** shows the associations between VLANs and interfaces. The **State** field shows whether the interfaces are up or down.

For the guest VLAN configuration, the interface is associated with the guest VLAN if and when the supplicant fails 802.1X user authentication.

Verifying That DHCP Snooping and IP Source Guard Are Working on the VLAN

Purpose Verify that DHCP snooping and IP source guard are enabled and working on the VLAN.

Action Send some DHCP requests from network devices (here they are DHCP clients) connected to the switch.

Use the **show dhcp snooping binding** command to display the DHCP snooping information when the interface on which the DHCP server connects to the switch is trusted. View the MAC addresses from which requests were sent and the IP addresses and leases provided by the server.

Use the **show ip-source-guard** command to view IP source guard information for the VLAN.

Meaning When the interface on which the DHCP server connects to the switch has been set to trusted, the output shows, for each MAC address, the assigned IP address and lease time—that is, the time, in seconds, remaining before the lease expires. Static IP addresses have no assigned lease time. Statically configured entries never expire.

The IP source guard database table contains the VLANs enabled for IP source guard, the untrusted access interfaces on those VLANs, the VLAN 802.1Q tag IDs if there are any, and the IP addresses and MAC addresses that are bound to one another. If a switch interface is associated with multiple VLANs and some of those VLANs are enabled for IP source guard and others are not, the VLANs that are not enabled for IP source guard have a star (*) in the **IP Address** and **MAC Address** fields.

Related Documentation

- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Example: Setting Up VoIP with 802.1X and LLDP-MED on a J-EX Series Switch on page 2302
- Example: Configuring IP Source Guard on a Data VLAN That Shares an Interface with a Voice VLAN on page 2608
- Configuring IP Source Guard (CLI Procedure) on page 2643

Example: Configuring IP Source Guard on a Data VLAN That Shares an Interface with a Voice VLAN

Ethernet LAN switches are vulnerable to attacks that involve spoofing (forging) of source IP addresses or source MAC addresses. These spoofed packets are sent from hosts connected to untrusted access interfaces on the switch. You can enable the IP source guard port security feature on J-EX Series switches to mitigate the effects of such attacks. If IP source guard determines that a source IP address and a source MAC address in a binding in an incoming packet are not valid, the switch does not forward the packet.

If two VLANs share an interface, you can configure IP source guard on just one of the VLANs; in this example, you configure IP source guard on an untagged data VLAN but not on the tagged voice VLAN. You can use 802.1X user authentication to validate the device connections on the data VLAN.

This example describes how to configure IP source guard with 802.1X user authentication on a data VLAN, with a voice VLAN on the same interface:

Requirements

This example uses the following hardware and software components:

- One J-EX Series switch
- A DHCP server to provide IP addresses to network devices on the switch
- A RADIUS server to provide 802.1X authentication

Before you configure IP source guard for the data VLANs, be sure you have:

- Connected the DHCP server to the switch.
- Connected the RADIUS server to the switch and configured user authentication on the server. See "Example: Connecting a RADIUS Server for 802.1X to a J-EX Series Switch" on page 2267.
- Configured the VLANs. See "Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches" on page 1070 for detailed information about configuring VLANs.

Overview and Topology

IP source guard checks the IP source address and MAC source address in a packet sent from a host attached to an untrusted access interface on the switch. If IP source guard determines that the packet header contains an invalid source IP address or source MAC address, it ensures that the switch does not forward the packet—that is, the packet is discarded.

When you configure IP source guard, you enable on it on one or more VLANs. IP source guard applies its checking rules to untrusted access interfaces on those VLANs. By default, on J-EX Series switches, access interfaces are untrusted and trunk interfaces are trusted. IP source guard does not check packets that have been sent to the switch by devices connected to either trunk interfaces or trusted access interfaces—that is, interfaces configured with **dhcp-trusted** so that a DHCP server can be connected to that interface to provide dynamic IP addresses.

IP source guard obtains information about IP-address/MAC-address/VLAN bindings from the DHCP snooping database. It causes the switch to validate incoming IP packets against the entries in that database.

The topology for this example includes one J-EX4200-24T switch, a PC and an IP phone connected on the same interface, a connection to a DHCP server, and a connection to a RADIUS server for user authentication.

NOTE: The 802.1X user authentication applied in this example is for single supplicants. Single-secure supplicant mode and multiple supplicant mode do not work with IP source guard. For more information about 802.1X authentication, see "Understanding Authentication on J-EX Series Switches" on page 2248.

TIP: You can set the ip-source-guard flag in the traceoptions statement for debugging purposes.

This example shows how to configure a static IP address to be added to the DHCP snooping database.

Configuration

CLI Quick Configuration

To quickly configure IP source guard on a data VLAN, copy the following commands and paste them into the switch terminal window:

```
set ethernet-switching-options voip interface ge-0/0/14.0 vlan voice
set ethernet-switching-options secure-access-port interface ge-0/0/24.0 dhcp-trusted
set ethernet-switching-options secure-access-port interface ge-0/0/14 static-ip 11.1.1.1 mac
00:11:11:11:11:11 vlan data
set ethernet-switching-options secure-access-port vlan data examine-dhcp
set ethernet-switching-options secure-access-port vlan data ip-source-guard
set interfaces ge-0/0/24 unit 0 family ethernet-switching vlan members data
set vlans voice vlan-id 100
set protocols lldp-med interface ge-0/0/14.0
set protocols dot1x authenticator authentication-profile-name profile52
set protocols dot1x authenticator interface ge-0/0/14.0 supplicant single
```

Step-by-Step Procedure

To configure IP source guard on the data VLAN:

1. Configure the VoIP interface:

```
[edit ethernet-switching-options]
user@switch# set voip interface ge-0/0/14.0 vlan voice
```

2. Configure the interface on which the DHCP server is connected to the switch as a trusted interface and add that interface to the data VLAN:

```
[edit ethernet-switching-options]
user@switch# set secure-access-port interface ge-0/0/24.0 dhcp-trusted
[edit interfaces]
user@switch# set ge-0/0/24 unit 0 family ethernet-switching vlan members data
```

3. Configure a static IP address on an interface on the data VLAN (optional)

```
[edit ethernet-switching-options]
user@switch# set secure-access-port interface ge-0/0/14 static-ip 11.1.1.1 mac
00:11:11:11:11:11 vlan data
```

4. Configure DHCP snooping and IP source guard on the data VLAN:

```
[edit ethernet-switching-options]
user@switch# set secure-access-port vlan data examine-dhcp
user@switch# set secure-access-port vlan data ip-source-guard
```

5. Configure 802.1X user authentication and LLDP-MED on the interface that is shared by the data VLAN and the voice VLAN:

```
[edit protocols]
user@switch# set lldp-med interface ge-0/0/14.0
user@switch# set dot1x authenticator authentication-profile-name profile52
user@switch# set dot1x authenticator interface ge-0/0/14.0 supplicant single
```

6. Set the VLAN ID for the voice VLAN:

```
[edit vlans]
user@switch# set voice vlan-id 100
```

Results Check the results of the configuration:

```
[edit ethernet-switching-options]
user@switch# show
voip {
  interface ge-0/0/14.0 {
    vlan voice;
  }
}
secure-access-port {
  interface ge-0/0/14.0 {
    static-ip 11.1.1.1 vlan data mac 00:11:11:11:11:11;
  }
  interface ge-0/0/24.0 {
    dhcp-trusted;
  }
  vlan data {
    examine-dhcp;
    ip-source-guard;
  }
}
```

```
[edit interfaces]
ge-0/0/24 {
  unit 0 {
    family ethernet-switching {
      vlan {
        members data;
      }
    }
  }
}
```

```
[edit vlans]
voice {
  vlan-id 100;
}
```

```
[edit protocols]
lldp-med {
  interface ge-0/0/14.0;
}
dot1x {
  authenticator {
```

```
        authentication-profile-name profile52;
        interface {
          ge-0/0/14.0 {
            supplicant single;
          }
        }
      }
    }
```

TIP: If you wanted to configure IP source guard on the voice VLAN as well as on the data VLAN, you would configure DHCP snooping and IP source guard exactly as you did for the data VLAN. The configuration result for the voice VLAN under secure-access-port would look like this:

```
secure-access-port {
  vlan voice {
    examine-dhcp;
    ip-source-guard;
  }
}
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying That 802.1X User Authentication Is Working on the Interface on page 2612
- Verifying the VLAN Association with the Interface on page 2613
- Verifying That DHCP Snooping and IP Source Guard Are Working on the Data VLAN on page 2613

Verifying That 802.1X User Authentication Is Working on the Interface

Purpose Verify the 802.1X configuration on interface **ge-0/0/14**.

Action Verify the 802.1X configuration with the operational mode command **show dot1x interface**:

```
user@switch> show dot1x interface ge-0/0/14.0 detail
ge-0/0/14.0
  Role: Authenticator
  Administrative state: Auto
  Supplicant mode: Single
  Number of retries: 3
  Quiet period: 60 seconds
  Transmit period: 30 seconds
  Mac Radius: Disabled
  Mac Radius Restrict: Disabled
  Reauthentication: Enabled
  Configured Reauthentication interval: 3600 seconds
  Supplicant timeout: 30 seconds
  Server timeout: 30 seconds
  Maximum EAPOL requests: 2
  Guest VLAN member: <not configured>
  Number of connected supplicants: 1
    Supplicant: user101, 00:04:0f:fd:ac:fe
```

```
Operational state: Authenticated
Authentication method: Radius
Authenticated VLAN: vo11
Dynamic Filter: <not configured>
Session Reauth interval: 60 seconds
Reauthentication due in 50 seconds
```

Meaning The **Supplicant mode** output field displays the configured administrative mode for each interface. Interface **ge-0/0/14.0** displays **Single** supplicant mode.

Verifying the VLAN Association with the Interface

Purpose Display the interface state and VLAN membership.

Action

```
user@switch>  show ethernet-switching interfaces
 Ethernet-switching table: 0 entries, 0 learned

user@switch> show ethernet-switching interfaces
Interface   State    VLAN members         Blocking
ge-0/0/0.0  down     default              unblocked
ge-0/0/1.0  down     employee             unblocked
ge-0/0/2.0  down     employee             unblocked
ge-0/0/12.0 down     default              unblocked
ge-0/0/13.0 down     default              unblocked
ge-0/0/13.0 down     vlan100              unblocked
ge-0/0/14.0 up       voice                unblocked
                     data                 unblocked
ge-0/0/17.0 down     employee             unblocked
ge-0/0/23.0 down     default              unblocked
ge-0/0/24.0 down     data                 unblocked
                     employee             unblocked
                     vlan100              unblocked
                     voice                unblocked
```

Meaning The field **VLAN members** shows that the **ge-0/0/14.0** interface supports both the **data** VLAN and the **voice** VLAN. The **State** field shows that the interface is up.

Verifying That DHCP Snooping and IP Source Guard Are Working on the Data VLAN

Purpose Verify that DHCP snooping and IP source guard are enabled and working on the data VLAN.

Action Send some DHCP requests from network devices (here they are DHCP clients) connected to the switch.

Display the DHCP snooping information when the interface on which the DHCP server connects to the switch is trusted. The following output results when requests are sent from the MAC addresses and the server has provided the IP addresses and leases:

```
user@switch>  show dhcp snooping binding
```

```
DHCP Snooping Information:
MAC address        IP address  Lease (seconds) Type     VLAN      Interface

00:05:85:3A:82:77  192.0.2.17  600             dynamic  employee  ge-0/0/1.0
00:05:85:3A:82:79  192.0.2.18  653             dynamic  employee  ge-0/0/1.0
00:05:85:3A:82:80  192.0.2.19  720             dynamic  employee  ge-0/0/2.0
00:05:85:3A:82:81  192.0.2.20  932             dynamic  employee  ge-0/0/2.0

                         00:30:48:92:A5:9D  10.10.10.7  720             dynamic
vlan100   ge-0/0/13.0
00:30:48:8D:01:3D  10.10.10.9  720             dynamic  data      ge-0/0/14.0
00:30:48:8D:01:5D  10.10.10.8  1230            dynamic  voice     ge-0/0/14.0
00:11:11:11:11:11  11.1.1.1    –               static   data      ge-0/0/14.0
00:05:85:27:32:88  192.0.2.22  –               static   employee  ge-0/0/17.0
00:05:85:27:32:89  192.0.2.23  –               static   employee  ge-0/0/17.0
00:05:85:27:32:90  192.0.2.27  –               static   employee  ge-0/0/17.0
```

View the IP source guard information for the data VLAN.

```
user@switch> show ip-source-guard
IP source guard information:
Interface    Tag  IP Address   MAC Address        VLAN

ge-0/0/13.0  0    10.10.10.7   00:30:48:92:A5:9D  vlan100

ge-0/0/14.0  0    10.10.10.9   00:30:48:8D:01:3D  data
ge-0/0/14.0  0    11.1.1.1     00:11:11:11:11:11  data

ge-0/0/13.0  100  *            *                  voice
```

Meaning When the interface on which the DHCP server connects to the switch has been set to trusted, the output (see the preceding sample output for **show dhcp snooping binding**) shows, for each MAC address, the assigned IP address and lease time—that is, the time, in seconds, remaining before the lease expires. Static IP addresses have no assigned lease time. Statically configured entries never expire.

The IP source guard database table contains the VLANs enabled for IP source guard, the untrusted access interfaces on those VLANs, the VLAN 802.1Q tag IDs if there are any, and the IP addresses and MAC addresses that are bound to one another. If a switch interface is associated with multiple VLANs and some of those VLANs are enabled for IP source guard and others are not, the VLANs that are not enabled for IP source guard have a star (*) in the **IP Address** and **MAC Address** fields. See the entry for the **voice** VLAN in the preceding sample output.

Related Documentation

- Example: Configuring IP Source Guard with Other J-EX Series Switch Features to Mitigate Address-Spoofing Attacks on Untrusted Access Interfaces on page 2600
- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Example: Setting Up VoIP with 802.1X and LLDP-MED on a J-EX Series Switch on page 2302
- Configuring IP Source Guard (CLI Procedure) on page 2643

Example: Setting Up DHCP Option 82 with a J-EX Series Switch as Relay Agent Between Clients and a DHCP Server

You can use DHCP option 82, also known as the DHCP relay agent information option, to help protect the J-EX Series switch against attacks such as spoofing (forging) of IP addresses and MAC addresses, and DHCP IP address starvation. Option 82 provides information about the network location of a DHCP client, and the DHCP server uses this information to implement IP addresses or other parameters for the client.

This example describes how to configure DHCP option 82 on a switch that is on the same VLAN with the DHCP clients but on a different VLAN from the DHCP server; the switch acts as a relay agent:

Requirements

This example uses the following hardware and software components:

- One J-EX4200-24T switch
- A DHCP server to provide IP addresses to network devices on the switch

Before you configure DHCP option 82 on the switch, be sure you have:

- Connected and configured the DHCP server.

  NOTE: Your DHCP server must be configured to accept DHCP option 82. If it is not configured for DHCP option 82, it does not use the DHCP option 82 information in the requests sent to it when it formulates its reply messages.

- Configured the **employee** VLAN on the switch and associated the interfaces on which the clients connect to the switch with that VLAN. See “Configuring VLANs for J-EX Series Switches (CLI Procedure)” on page 1136.
- Configured the **corporate** VLAN for the DHCP server.
- Configured the switch as a BOOTP relay agent. See “DHCP/BOOTP Relay for J-EX Series Switches Overview” on page 446.
- Configured the routed VLAN interface (RVI) to allow the switch to relay packets to the server and receive packets from the server. See “Configuring Routed VLAN Interfaces (CLI Procedure)” on page 1137.

Overview and Topology

If DHCP option 82 is enabled on the switch, then when a network device—a DHCP client—that is connected to the switch on an untrusted interface sends a DHCP request, the switch inserts information about the client's network location into the packet header

of that request. The switch then sends the request (in this setting, it relays the request) to the DHCP server. The DHCP server reads the option 82 information in the packet header and uses it to implement the IP address or other parameter for the client.

When option 82 is enabled on the switch, then this sequence of events occurs when a DHCP client sends a DHCP request:

1. The switch receives the request and inserts the option 82 information in the packet header.
2. The switch relays the request to the DHCP server.
3. The server uses the DHCP option 82 information to formulate its reply and sends a response back to the switch. It does not alter the option 82 information.
4. The switch strips the option 82 information from the response packet.
5. The switch forwards the response packet to the client.

In this example, you configure option 82 on the J-EX Series switch. The switch is configured as a BOOTP relay agent. The switch connects to the DHCP server through the routed VLAN interface (RVI) that you configured. The switch and clients are members of the **employee** VLAN. The DHCP server is a member of the **corporate** VLAN.

Configuration

To configure DHCP option 82:

CLI Quick Configuration

To quickly configure DHCP option 82, copy the following commands and paste them into the switch terminal window:

```
set forwarding-options helpers bootp dhcp-option82
set forwarding-options helpers bootp dhcp-option82 circuit-id prefix hostname
set forwarding-options helpers bootp dhcp-option82 circuit-id use-vlan-id
set forwarding-options helpers bootp dhcp-option82 remote-id
set forwarding-options helpers bootp dhcp-option82 remote-id prefix mac
set forwarding-options helpers bootp dhcp-option82 remote-id use-string employee-switch1
set forwarding-options helpers bootp dhcp-option82 vendor-id
```

Step-by-Step Procedure

To configure DHCP option 82:

1. Specify DHCP option 82 for the **employee** VLAN:

   ```
   [edit forwarding-options helpers bootp]
   user@switch# set dhcp-option82
   ```

2. Configure a prefix for the circuit ID suboption (the prefix is always the hostname of the switch):

   ```
   [edit forwarding-options helpers bootp]
   user@switch# set dhcp-option82 circuit-id prefix hostname
   ```

3. Specify that the circuit ID suboption value contains the VLAN ID rather than the VLAN name (the default):

   ```
   [edit forwarding-options helpers bootp]
   user@switch# set dhcp-option82 circuit-id use-vlan-id
   ```

4. Specify that the remote ID suboption be included in the DHCP option 82 information:

```
[edit forwarding-options helpers bootp]
user@switch# set dhcp-option82 remote-id
```

5. Configure a prefix for the remote ID suboption (here, the prefix is the MAC address of the switch):

```
[edit forwarding-options helpers bootp]
user@switch# set dhcp-option82 remote-id prefix mac
```

6. Specify that the remote ID suboption value contains a character string (here, the string is **employee-switch1**):

```
[edit forwarding-options helpers bootp]
user@switch# set dhcp-option82 remote-id use-string employee-switch1
```

7. Configure a vendor ID suboption value, and use the default value. To use the default value, do not type a character string after the **vendor-id** option keyword:

```
[edit forwarding-options helpers bootp]
user@switch# set dhcp-option82 vendor-id
```

Results Check the results of the configuration:

```
[edit forwarding-options helpers bootp]
user@switch# show

  dhcp-option82 {
    circuit-id {
      prefix hostname;
      use-vlan-id;
    }
    remote-id {
      prefix mac;
      use-string employee-switch1;
    }
    vendor-id;
  }
```

Related Documentation

- Example: Setting Up DHCP Option 82 on a J-EX Series Switch with No Relay Agent Between Clients and DHCP Server on page 2617
- Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 2646
- RFC 3046, *DHCP Relay Agent Information Option*, at http://tools.ietf.org/html/rfc3046.

Example: Setting Up DHCP Option 82 on a J-EX Series Switch with No Relay Agent Between Clients and DHCP Server

You can use DHCP option 82, also known as the DHCP relay agent information option, to help protect the J-EX Series switch against attacks such as spoofing (forging) of IP addresses and MAC addresses, and DHCP IP address starvation. Option 82 provides information about the network location of a DHCP client, and the DHCP server uses this information to implement IP addresses or other parameters for the client.

This example describes how to configure DHCP option 82 on a switch with DHCP clients, DHCP server, and switch all on the same VLAN:

Requirements

This example uses the following hardware and software components:

- One J-EX Series switch
- A DHCP server to provide IP addresses to network devices on the switch

Before you configure DHCP option 82 on the switch, be sure you have:

- Connected and configured the DHCP server.

NOTE: Your DHCP server must be configured to accept DHCP option 82. If it is not configured for DHCP option 82, it does not use the DHCP option 82 information in the requests sent to it when it formulates its reply messages.

- Configured the **employee** VLAN on the switch and associated the interfaces on which the clients and the server connect to the switch with that VLAN. See “Configuring VLANs for J-EX Series Switches (CLI Procedure)” on page 1136.

Overview and Topology

If DHCP option 82 is enabled on the switch, then when a network device—a DHCP client—that is connected to the switch on an untrusted interface sends a DHCP request, the switch inserts information about the client's network location into the packet header of that request. The switch then sends the request to the DHCP server. The DHCP server reads the option 82 information in the packet header and uses it to implement the IP address or other parameter for the client.

DHCP option 82 is enabled on an individual VLAN or on all VLANs on the switch.

When option 82 is enabled on the switch, then this sequence of events occurs when a DHCP client sends a DHCP request:

1. The switch receives the request and inserts the option 82 information in the packet header.
2. The switch forwards the request to the DHCP server.
3. The server uses the DHCP option 82 information to formulate its reply and sends a response back to the switch. It does not alter the option 82 information.
4. The switch strips the option 82 information from the response packet.
5. The switch forwards the response packet to the client.

Figure 74 on page 2619 illustrates the topology for this example.

Figure 74: Network Topology for Configuring DHCP Option 82 on a Switch That Is on the Same VLAN as the DHCP Clients and the DHCP Server

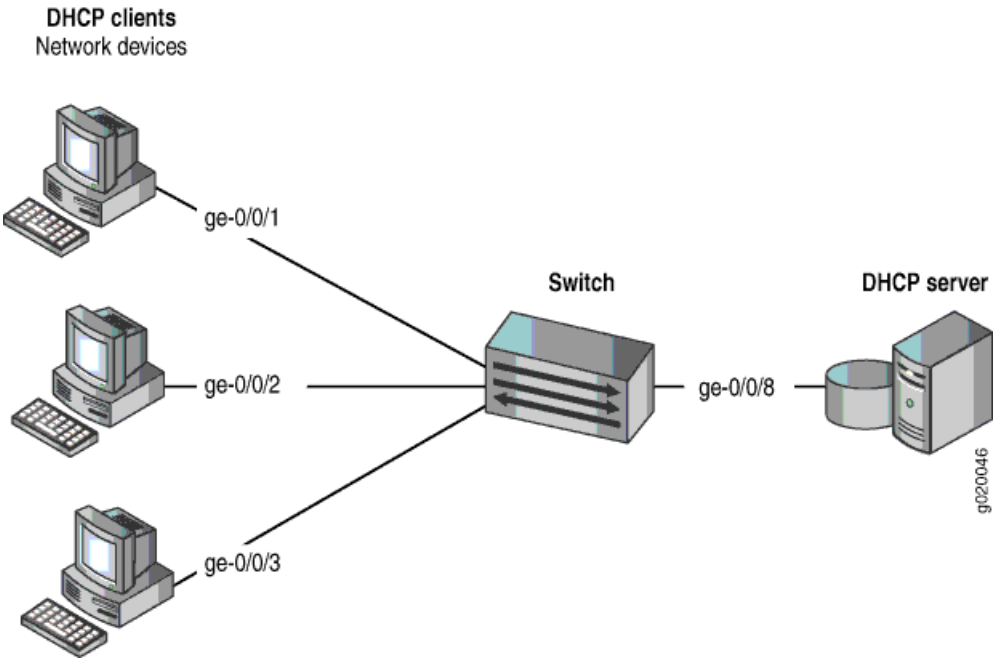

In this example, you configure DHCP option 82 on the J-EX Series switch. The switch connects to the DHCP server on interface **ge-0/0/8**. The DHCP clients connect to the switch on interfaces **ge-0/0/1**, **ge-0/0/2**, and **ge-0/0/3**. The switch, server, and clients are all members of the **employee** VLAN.

Configuration

To configure DHCP option 82:

CLI Quick Configuration

To quickly configure DHCP option 82, copy the following commands and paste them into the switch terminal window:

```
set ethernet-switching-options secure-access-port vlan employee dhcp-option82
set ethernet-switching-options secure-access-port vlan employee dhcp-option82 circuit-id prefix hostname
set ethernet-switching-options secure-access-port vlan employee dhcp-option82 circuit-id use-vlan-id
set ethernet-switching-options secure-access-port vlan employee dhcp-option82 remote-id
set ethernet-switching-options secure-access-port vlan employee dhcp-option82 remote-id prefix mac
set ethernet-switching-options secure-access-port vlan employee dhcp-option82 remote-id use-string employee-switch1
set ethernet-switching-options secure-access-port vlan employee dhcp-option82 vendor-id
```

Step-by-Step Procedure

To configure DHCP option 82:

1. Specify DHCP option 82 for the **employee** VLAN:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee dhcp-option82
```

2. Configure a prefix for the circuit ID suboption (the prefix is always the hostname of the switch):

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee dhcp-option82 circuit-id prefix hostname
```

3. Specify that the circuit ID suboption value contains the VLAN ID rather than the VLAN name (the default):

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee dhcp-option82 circuit-id use-vlan-id
```

4. Specify that the remote ID suboption be included in the DHCP option 82 information:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee dhcp-option82 remote-id
```

5. Configure a prefix for the remote ID suboption (here, the prefix is the MAC address of the switch):

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee dhcp-option82 remote-id prefix mac
```

6. Specify that the remote ID suboption value contains a character string (here, the string is **employee-switch1**):

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee dhcp-option82 remote-id use-string employee-switch1
```

7. Configure a vendor ID suboption value, and use the default value. To use the default value, do not type a character string after the **vendor-id** option keyword:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee dhcp-option82 vendor-id
```

Results

Check the results of the configuration:

```
[edit ethernet-switching-options secure-access-port]
user@switch# show

  vlan employee {
    dhcp-option82 {
      circuit-id {
        prefix hostname;
        use-vlan-id;
      }
      remote-id {
        prefix mac;
        use-string employee-switch1;
      }
      vendor-id;
    }
  }
```

Related Documentation

- Example: Setting Up DHCP Option 82 with a J-EX Series Switch as Relay Agent Between Clients and a DHCP Server on page 2615
- Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 2649
- RFC 3046, *DHCP Relay Agent Information Option*, at http://tools.ietf.org/html/rfc3046.

Example: Configuring Proxy ARP on a J-EX Series Switch

You can configure proxy Address Resolution Protocol (ARP) on your J-EX Series switch to enable the switch to respond to ARP queries for network addresses by offering its own MAC address. With proxy ARP enabled, the switch captures and routes traffic to the intended destination.

This example shows how to configure proxy ARP on an access switch:

Requirements

This example uses the following hardware and software components:

- One J-EX Series switch

Overview and Topology

This example shows the configuration of proxy ARP on an interface of a J-EX Series switch using restricted mode. In restricted mode, the switch does not proxy for hosts on the same subnet.

The topology for this example consists of one J-EX Series switch. When a host wants to communicate with a host that is not already in its ARP table, it broadcasts an ARP request for the MAC address of the destination host:

- When proxy ARP is not enabled, a host that shares the same IP address replies directly to the ARP request, providing its MAC address, and future transmissions are sent directly to the destination host MAC address.
- When proxy ARP is enabled, the switch responds to ARP requests, providing the switch's MAC address—even when the destination IP address is the same as the source IP address. Thus, communications must be sent through the switch and then routed through the switch to the appropriate destination.

Configuration

To configure proxy ARP, perform the following tasks:

CLI Quick Configuration

To quickly configure proxy ARP on an interface, copy the following command and paste it into the switch terminal window:

```
[edit]
set interfaces ge-0/0/3 unit 0 proxy-arp restricted
```

Step-by-Step Procedure

You configure proxy ARP on individual interfaces.

1. To configure proxy ARP on an interface:

```
[edit interfaces]
user@switch# set ge-0/0/3 unit 0 proxy-arp restricted
```

BEST PRACTICE: We recommend that you configure proxy ARP in restricted mode. In restricted mode, the switch does not act as proxy if the source and target IP addresses are on the same subnet. If you use unrestricted mode, disable gratuitous ARP requests on the interface to avoid the situation of the switch's response to a gratuitous ARP request appearing to the host to be an indication of an IP conflict:

```
[edit interfaces]
user@switch# set ge-0/0/3 no-gratuitous-arp-request
```

Results

Display the results of the configuration:

```
user@switch> show configuration
  interfaces {
    ge-0/0/3 {
      unit 0 {
        proxy-arp restricted;
        family ethernet-switching;
      }
    }
```

Verification

To verify that the switch is sending proxy ARP messages, perform these tasks:

- Verifying That the Switch Is Sending Proxy ARP Messages on page 2622

Verifying That the Switch Is Sending Proxy ARP Messages

Purpose

Verify that the switch is sending proxy ARP messages.

Action

List the system statistics for ARP messages:

```
user@switch> show system statistics arp
arp:
        198319 datagrams received
        45 ARP requests received
        12 ARP replies received
        2 resolution requests received
        2 unrestricted proxy requests
        0 restricted proxy requests
        0 received proxy requests
```

```
0 proxy requests not proxied
0 restricted-proxy requests not proxied
0 with bogus interface
0 with incorrect length
0 for non-IP protocol
0 with unsupported op code
0 with bad protocol address length
0 with bad hardware address length
0 with multicast source address
0 with multicast target address
0 with my own hardware address
168705 for an address not on the interface
0 with a broadcast source address
0 with source address duplicate to mine
29555 which were not for me
0 packets discarded waiting for resolution
4 packets sent after waiting for resolution
27 ARP requests sent
47 ARP replies sent
0 requests for memory denied
0 requests dropped on entry
0 requests dropped during retry
0 requests dropped due to interface deletion
0 requests on unnumbered interfaces
0 new requests on unnumbered interfaces
0 replies for from unnumbered interfaces
0 requests on unnumbered interface with non-subnetted donor
0 replies from unnumbered interface with non-subnetted donor
```

Meaning The statistics show that two proxy ARP requests were received, and the **proxy requests not proxied** field indicates that all the unproxied ARP requests received have been proxied by the switch.

Related Documentation

- Configuring Proxy ARP (CLI Procedure) on page 1153
- Understanding Proxy ARP on J-EX Series Switches on page 1059

CHAPTER 95

Configuring Port Security

Configuring Port Security (CLI Procedure)

Ethernet LANs are vulnerable to attacks such as address spoofing and Layer 2 denial of service (DoS) on network devices. Port security features such as DHCP snooping, DAI (dynamic ARP inspection), MAC limiting, and MAC move limiting, as well as trusted DHCP server, help protect the access ports on your J-EX Series switch against the losses of information and productivity that can result from such attacks.

To configure port security features using the CLI:

1. Enable DHCP snooping:
   - On a specific VLAN:

     ```
     [edit ethernet-switching-options secure-access port]
     user@switch# set vlan default examine-dhcp
     ```

   - On all VLANs:

     ```
     [edit ethernet-switching-options secure-access port]
     user@switch# set vlan all examine-dhcp
     ```

2. Enable DAI:
   - On a single VLAN (here, the VLAN is **employee-vlan**):

     ```
     [edit ethernet-switching-options secure-access-port]
     user@switch# set vlan employee–vlan arp-inspection
     ```

   - On all VLANs:

     ```
     [edit ethernet-switching-options secure-access-port]
     user@switch# set vlan all arp-inspection
     ```

3. Limit the number of dynamic MAC addresses and specify the action to take if the limit is exceeded—for example, set a MAC limit of **5** with an action of **drop**:
   - On a single interface (here, the interface is **ge-0/0/1**):

     ```
     [edit ethernet-switching-options secure-access-port]
     user@switch# set interface ge–0/0/1 mac-limit 5 action drop
     ```

   - On all interfaces:

     ```
     [edit ethernet-switching-options secure-access-port]
     user@switch# set interface all mac–limit 5 action drop
     ```

4. Specify allowed MAC addresses:
   - On a single interface (here, the interface is **ge-0/0/2**):

     ```
     [edit ethernet-switching-options secure-access-port]
     user@switch# set interface ge–0/0/2 allowed-mac 00:05:85:3A:82:80
     user@switch# set interface ge–0/0/2 allowed-mac 00:05:85:3A:82:81
     user@switch# set interface ge–0/0/2 allowed-mac 00:05:85:3A:82:83
     ```

   - On all interfaces:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface all allowed-mac 00:05:85:3A:82:80
user@switch# set interface all allowed-mac 00:05:85:3A:82:81
user@switch# set interface all allowed-mac 00:05:85:3A:82:83
```

5. Limit the number of times a MAC address can move from its original interface in one second—for example, set a MAC move limit of **5** with an action of **drop** if the limit is exceeded:

   - On a single VLAN (here, the VLAN is **employee-vlan**):

     ```
     [edit ethernet-switching-options secure-access-port]
     user@switch# set vlan employee–vlan mac-move-limit 5 action drop
     ```

   - On all VLANs:

     ```
     [edit ethernet-switching-options secure-access-port]
     user@switch# set vlan all mac–move-limit 5 action drop
     ```

6. Configure a trusted DHCP server on an interface (here, the interface is **ge-0/0/8**):

   ```
   [edit ethernet-switching-options secure-access port]
   user@switch# set interface ge–0/0/8 dhcp-trusted
   ```

Related Documentation

- Configuring Port Security (J-Web Procedure) on page 2627
- Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure) on page 2516
- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Example: Configuring DHCP Snooping, DAI , and MAC Limiting on a J-EX Series Switch with Access to a DHCP Server Through a Second Switch on page 2593
- Monitoring Port Security on page 2653
- Port Security for J-EX Series Switches Overview on page 2545

Configuring Port Security (J-Web Procedure)

To configure port security on a J-EX Series switch using the J-Web interface:

1. Select **Configure** > **Security** > **Port Security**.

   The **VLAN List** table lists all the VLAN names, VLAN identifiers, port members, and port security VLAN features.

   The **Interface List** table lists all the ports and indicates whether security features have been enabled on the ports.

NOTE: After you make changes to the configuration in this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See "Using the Commit Options to Commit Configuration Changes (J-Web Procedure)" on page 334 for details about all commit options.

2. Click one:

   - **Edit**—Click this option to modify the security features for the selected port or VLAN.

     Enter information as specified in Table 328 on page 2628 to modify Port Security settings on VLANs.

     Enter information as specified in Table 329 on page 2629 to modify Port Security settings on interfaces.

   - **Activate/Deactivate**—Click this option to enable or disable security on the switch.

Table 328: Port Security Settings on VLANs

Field	Function	Your Action
Enable DHCP Snooping on VLAN	Allows the switch to monitor and control DHCP messages received from untrusted devices connected to the switch. Builds and maintains a database of valid IP addresses/MAC address bindings. (By default, access ports are untrusted and trunk ports are trusted.)	Select to enable DHCP snooping on a specified VLAN or all VLANs. TIP: For private VLANs (PVLANs), enable DHCP snooping on the primary VLAN. If you enable DHCP snooping only on a community VLAN, DHCP messages coming from PVLAN trunk ports are not snooped.
Enable ARP Inspection on VLAN	Uses information in the DHCP snooping database to validate ARP packets on the LAN and protect against ARP cache poisoning.	Select to enable ARP inspection on a specified VLAN or all VLANs. (Configure any port on which you do not want ARP inspection to occur as a trusted DHCP server port.)
MAC Movement	Specifies the number of times per second that a MAC address can move to a new interface.	Enter a number. The default is unlimited.
MAC Movement Action	Specifies the action to be taken if the MAC move limit is exceeded.	Select one: • Log—Generate a system log entry, an SNMP trap, or an alarm. • Drop—Drop the packets and generate a system log entry, an SNMP trap, or an alarm (default). • Shutdown—Shut down the VLAN and generate an alarm. You can mitigate the effect of this option by configuring autorecovery from the disabled state and specifying a **disable timeout** value. See "Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure)" on page 2516. • None—No action to be taken.

Table 329: Port Security on Interfaces

Field	Function	Your Action
Trust DHCP	Specifies trusting DHCP packets on the selected interface. By default, trunk ports are **dhcp-trusted**.	Select to enable DHCP trust.
MAC Limit	Specifies the number of MAC addresses that can be learned on a single Layer 2 access port. This option is not valid for trunk ports.	Enter a number.
MAC Limit Action	Specifies the action to be taken if the MAC limit is exceeded. This option is not valid for trunk ports.	Select one: • Log—Generate a system log entry, an SNMP trap, or an alarm. • Drop—Drop the packets and generate a system log entry, an SNMP trap, or an alarm. (Default) • Shutdown—Shut down the interface and generate an alarm. You can mitigate the effect of this option by configuring autorecovery from the disabled state and specifying a **disable timeout** value. See "Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure)." on page 2516 • None—No action to be taken.
Allowed MAC List	Specifies the MAC addresses that are allowed for the interface.	To add a MAC address: 1. Click **Add**. 2. Enter the MAC address. 3. Click **OK**.

Related Documentation

- Configuring Port Security (CLI Procedure) on page 2626
- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Monitoring Port Security on page 2653
- Port Security for J-EX Series Switches Overview on page 2545

Enabling DHCP Snooping (CLI Procedure)

DHCP snooping allows the switch to monitor and control DHCP messages received from untrusted devices connected to the J-EX Series switch. It builds and maintains a database of valid IP-address/MAC-address (IP-MAC) bindings called the DHCP snooping database.

You configure DHCP snooping for each VLAN, not for each interface (port). By default, DHCP snooping is disabled for all VLANs.

To enable DHCP snooping on a VLAN or all VLANs by using the CLI:

- On a specific VLAN (here, the VLAN is **default**):

```
[edit ethernet-switching-options secure-access port]
user@switch# set vlan default examine-dhcp
```

- On all VLANs:

```
[edit ethernet-switching-options secure-access port]
user@switch# set vlan all examine-dhcp
```

TIP: By default, the IP-MAC bindings are lost when the switch is rebooted and DHCP clients (the network devices, or hosts) must reacquire bindings. However, you can configure the bindings to persist by setting the dhcp-snooping-file statement to store the database file either locally or remotely.

TIP: For private VLANs (PVLANs), enable DHCP snooping on the primary VLAN. If you enable DHCP snooping only on a community VLAN, DHCP messages coming from PVLAN trunk ports are not snooped.

Related Documentation

- Enabling DHCP Snooping (J-Web Procedure) on page 2631
- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Example: Configuring DHCP Snooping, DAI , and MAC Limiting on a J-EX Series Switch with Access to a DHCP Server Through a Second Switch on page 2593
- Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks on page 2586
- Verifying That DHCP Snooping Is Working Correctly on page 2654
- Monitoring Port Security on page 2653
- Understanding DHCP Snooping for Port Security on J-EX Series Switches on page 2549

Enabling DHCP Snooping (J-Web Procedure)

DHCP snooping allows the J-EX Series switch to monitor and control DHCP messages received from untrusted devices connected to the switch. It builds and maintains a database of valid IP-address/MAC-address (IP-MAC) bindings called the DHCP snooping database.

You configure DHCP snooping for each VLAN, not for each interface (port). By default, DHCP snooping is disabled for all VLANs.

To enable DHCP snooping on one or more VLANs by using the J-Web interface:

1. Select **Configure>Security>Port Security**.
2. Select one or more VLANs from the VLAN list.
3. Click the **Edit** button. If a message appears asking if you want to enable port security, click **Yes**.
4. Select the **Enable DHCP Snooping on VLAN** check box and then click **OK**.
5. Click **OK** after the command has been successfully delivered.

NOTE: You can enable or disable port security on the switch at any time by clicking the **Activate** or **Deactivate** button on the Port Security Configuration page. If security status is shown as **Disabled** when you try to edit settings for any VLANs or interfaces (ports), the message asking if you want to enable port security appears.

Related Documentation

- Enabling DHCP Snooping (CLI Procedure) on page 2630
- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Example: Configuring DHCP Snooping, DAI , and MAC Limiting on a J-EX Series Switch with Access to a DHCP Server Through a Second Switch on page 2593
- Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks on page 2586
- Verifying That DHCP Snooping Is Working Correctly on page 2654
- Monitoring Port Security on page 2653
- Understanding DHCP Snooping for Port Security on J-EX Series Switches on page 2549

Enabling a Trusted DHCP Server (CLI Procedure)

You can configure any interface on the J-EX Series switch that connects to a DHCP server as a trusted interface (port). Configuring a DHCP server on a trusted interface protects against rogue DHCP servers sending leases.

You configure a trusted DHCP server on an interface, not on a VLAN. By default, all access interfaces are untrusted and all trunk interfaces are trusted.

To configure a trusted interface for a DHCP server by using the CLI (here, the interface is **ge-0/0/8**):

```
[edit ethernet-switching-options secure-access port]
user@switch# set interface ge–0/0/8 dhcp-trusted
```

Related Documentation

- Enabling a Trusted DHCP Server (J-Web Procedure) on page 2632
- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Example: Configuring a DHCP Server Interface as Untrusted to Protect the Switch from Rogue DHCP Server Attacks on page 2579
- Verifying That a Trusted DHCP Server Is Working Correctly on page 2655
- Monitoring Port Security on page 2653
- Understanding Trusted DHCP Servers for Port Security on J-EX Series Switches on page 2559

Enabling a Trusted DHCP Server (J-Web Procedure)

You can configure any interface on the J-EX Series switch that connects to a DHCP server as a trusted interface (port). Configuring a DHCP server on a trusted interface protects against rogue DHCP servers sending leases.

You configure a trusted DHCP server on an interface, not on a VLAN. By default, all access interfaces are untrusted and all trunk interfaces are trusted.

To enable a trusted DHCP server on one or more interfaces by using the J-Web interface:

1. Select **Configure>Security>Port Security**.
2. Select one or more interfaces from the Port list.
3. Click the **Edit** button. If a message appears asking if you want to enable port security, click **Yes**.
4. Select the **Trust DHCP** check box and then click **OK**.
5. Click **OK** after the command has been successfully delivered.

NOTE: You can enable or disable port security on the switch at any time by clicking the **Activate** or **Deactivate** button on the Port Security Configuration page. If security status is shown as **Disabled** when you try to edit settings for any VLANs or interfaces (ports), the message asking if you want to enable port security appears.

Related Documentation

- Enabling a Trusted DHCP Server (CLI Procedure) on page 2632
- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Example: Configuring a DHCP Server Interface as Untrusted to Protect the Switch from Rogue DHCP Server Attacks on page 2579
- Verifying That a Trusted DHCP Server Is Working Correctly on page 2655
- Monitoring Port Security on page 2653
- Understanding Trusted DHCP Servers for Port Security on J-EX Series Switches on page 2559

Enabling Dynamic ARP Inspection (CLI Procedure)

Dynamic ARP inspection (DAI) protects J-EX Series switches against ARP spoofing. DAI inspects ARP packets on the LAN and uses the information in the DHCP snooping database on the switch to validate ARP packets and to protect against ARP cache poisoning.

You configure DAI for each VLAN, not for each interface (port). By default, DAI is disabled for all VLANs.

To enable dynamic ARP inspection (DAI) on a VLAN or all VLANs using the CLI:

- On a single VLAN (here, the VLAN is **employee-vlan**):

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee-vlan arp-inspection
```

- On all VLANs:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan all arp-inspection
```

Related Documentation

- Enabling Dynamic ARP Inspection (J-Web Procedure) on page 2634
- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Example: Configuring DHCP Snooping, DAI , and MAC Limiting on a J-EX Series Switch with Access to a DHCP Server Through a Second Switch on page 2593

- Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks on page 2586
- Verifying That DAI Is Working Correctly on page 2656
- Monitoring Port Security on page 2653
- Understanding DAI for Port Security on J-EX Series Switches on page 2555

Enabling Dynamic ARP Inspection (J-Web Procedure)

Dynamic ARP inspection (DAI) protects J-EX Series switches against ARP spoofing. DAI inspects ARP packets on the LAN and uses the information in the DHCP snooping database on the switch to validate ARP packets and to protect against ARP cache poisoning.

You configure DAI for each VLAN, not for each interface (port). By default, DAI is disabled for all VLANs.

To enable DAI on one or more VLANs by using the J-Web interface:

1. Select **Configure>Security>Port Security**.
2. Select one or more VLANs from the VLAN list.
3. Click the **Edit** button. If a message appears asking if you want to enable port security, click **Yes**.
4. Select the **Enable ARP Inspection on VLAN** check box and then click **OK**.
5. Click **OK** after the command has been successfully delivered.

NOTE: You can enable or disable port security on the switch at any time by clicking the Activate or Deactivate button on the Port Security Configuration page. If security status is shown as Disabled when you try to edit settings for any VLANs or interfaces (ports), the message asking if you want to enable port security appears.

Related Documentation

- Enabling Dynamic ARP Inspection (CLI Procedure) on page 2633
- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Example: Configuring DHCP Snooping, DAI , and MAC Limiting on a J-EX Series Switch with Access to a DHCP Server Through a Second Switch on page 2593
- Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks on page 2586
- Verifying That DAI Is Working Correctly on page 2656
- Monitoring Port Security on page 2653

- Understanding DAI for Port Security on J-EX Series Switches on page 2555

Configuring MAC Limiting (CLI Procedure)

MAC limiting protects against flooding of the Ethernet switching table on the J-EX Series switch. MAC limiting sets a limit on the number of MAC addresses that can be learned on a single Layer 2 access interface (port).

The Junos OS provides two MAC limiting methods:

- Maximum number of dynamic MAC addresses allowed per interface—When the limit is exceeded, incoming packets with new MAC addresses are dropped.
- Specific "allowed" MAC addresses for the access interface—Any MAC address that is not in the list of configured addresses is not learned and the switch logs the message.

NOTE: If you do not want the switch to log messages received for invalid MAC addresses on an interface that has been configured for specific "allowed" MAC addresses, you can disable the logging by configuring the **no-allowed-mac-log** statement.

You configure MAC limiting per interface, not per VLAN. You can specify the maximum number of dynamic MAC addresses that can be learned on a single Layer 2 access interface or on all Layer 2 access interfaces.

You can choose to have one of the following actions performed when the limit of MAC addresses is exceeded:

- **drop**—Drop the packet and generate an alarm, an SNMP trap, or a system log entry. This is the default.
- **log**—Do not drop the packet but generate an alarm, an SNMP trap, or a system log entry.
- **none**—Take no action.
- **shutdown**—Disable the interface and generate an alarm. If you have configured the switch with the **port-error-disable** statement, the disabled interface recovers automatically upon expiration of the specified disable timeout. If you have not configured the switch for autorecovery from port error disabled conditions, you can bring up the disabled interfaces by running the **clear ethernet-switching port-error** command.

To configure MAC limiting on a specific interface or on all interfaces, using the CLI:

1. For limiting the number of dynamic MAC addresses, set a MAC limit of **5**.

   The action is not specified, so the switch performs the default action **drop** if the limit is exceeded:

   - On a single interface (here, the interface is **ge-0/0/1**):

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface ge–0/0/1 mac-limit 5
```

- On all interfaces:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface all mac–limit 5
```

NOTE: Do not set the mac-limit to 1. The first learned MAC address is often inserted into the forwarding database automatically (for instance, for Routed VLAN Interfaces the first MAC address inserted into the forwarding database is the MAC address of the RVI. For Aggregated Ethernet bundles using LACP, the first MAC address inserted into the forwarding database in the forwarding table is the source address of the protocol packet). The switch will therefore not learn MAC addresses other than the automatic addresses when the mac-limit is set to 1, and this will cause problems with MAC learning and forwarding.

2. For specifying specific allowed MAC addresses:

   - On a single interface (here, the interface is **ge-0/0/2**):

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface ge–0/0/2 allowed-mac 00:05:85:3A:82:80
user@switch# set interface ge–0/0/2 allowed-mac 00:05:85:3A:82:81
user@switch# set interface ge–0/0/2 allowed-mac 00:05:85:3A:82:83
```

   - On all interfaces:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface all allowed-mac 00:05:85:3A:82:80
user@switch# set interface all allowed-mac 00:05:85:3A:82:81
user@switch# set interface all allowed-mac 00:05:85:3A:82:83
```

Related Documentation

- Configuring MAC Limiting (J-Web Procedure) on page 2637
- Example: Configuring MAC Limiting, Including Dynamic and Allowed MAC Addresses, to Protect the Switch from Ethernet Switching Table Overflow Attacks on page 2576
- Verifying That MAC Limiting Is Working Correctly on page 2657
- Setting the none Action on an Interface to Override a MAC Limit Applied to All Interfaces (CLI Procedure) on page 2642
- Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure) on page 2516
- Understanding MAC Limiting and MAC Move Limiting for Port Security on J-EX Series Switches on page 2557
- **no-allowed-mac-log on page 2686**

Configuring MAC Limiting (J-Web Procedure)

MAC limiting protects against flooding of the Ethernet switching table on a J-EX Series switch. MAC limiting sets a limit on the number of MAC addresses that can be learned on a single Layer 2 access interface (port).

The Junos OS provides two MAC limiting methods:

- Maximum number of dynamic MAC addresses allowed per interface—If the limit is exceeded, incoming packets with new MAC addresses are dropped.
- Specific "allowed" MAC addresses for the access interface—Any MAC address that is not in the list of configured addresses is not learned.

You configure MAC limiting for each interface, not for each VLAN. You can specify the maximum number of dynamic MAC addresses that can be learned on a single Layer 2 access interface or on all Layer 2 access interfaces. The default action that the switch will take if that maximum number is exceeded is **drop**—drop the packet and generate an alarm, an SNMP trap, or a system log entry.

To enable MAC limiting on one or more interfaces using the J-Web interface:

1. Select **Configure>Security>Port Security**.
2. Select one or more interfaces from the **Interface List**.
3. Click the **Edit** button. If a message appears asking whether you want to enable port security, click **Yes**.
4. To set a dynamic MAC limit:
    1. Type a limit value in the **MAC Limit** box.
    2. Select an action from the **MAC Limit Action** box (optional). The switch takes this action when the MAC limit is exceeded. If you do not select an action, the switch applies the default action, **drop**.
        - Log—Generate a system log entry, an SNMP trap, or an alarm.
        - Drop—Drop the packets and generate a system log entry, an SNMP trap, or an alarm. (Default)
        - Shutdown—Shut down the VLAN and generate an alarm. You can mitigate the effect of this option by configuring the switch for autorecovery from the disabled state and specifying a **disable timeout** value. See "Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure)" on page 2516. If you have not configured autorecovery from the disabled state, you can bring up the interfaces by running the **clear ethernet-switching port-error** command.
        - None— No action to be taken.
5. To add allowed MAC addresses:

1. Click **Add**.
2. Type the allowed MAC address and click **OK**.

   Repeat this step to add more allowed MAC addresses.

6. Click **OK** when you have finished setting MAC limits.
7. Click **OK** after the configuration has been successfully delivered.

NOTE: You can enable or disable port security on the switch at any time by clicking the **Activate** or **Deactivate** button on the Port Security Configuration page. If security status is shown as **Disabled** when you try to edit settings for any VLANs or interfaces (ports), a message asking whether you want to enable port security appears.

Related Documentation

- Configuring MAC Limiting (CLI Procedure) on page 2635
- Example: Configuring Allowed MAC Addresses to Protect the Switch from DHCP Snooping Database Alteration Attacks on page 2590
- Example: Configuring MAC Limiting, Including Dynamic and Allowed MAC Addresses, to Protect the Switch from Ethernet Switching Table Overflow Attacks on page 2576
- Example: Configuring MAC Limiting to Protect the Switch from DHCP Starvation Attacks on page 2583
- Verifying That MAC Limiting Is Working Correctly on page 2657
- Setting the none Action on an Interface to Override a MAC Limit Applied to All Interfaces (CLI Procedure) on page 2642
- Understanding MAC Limiting and MAC Move Limiting for Port Security on J-EX Series Switches on page 2557

Configuring MAC Move Limiting (CLI Procedure)

MAC move limiting detects MAC address movement and MAC address spoofing on access ports. MAC address movements are tracked, and if a MAC address moves more than the configured number of times within one second, the configured (or default) action is performed. You enable this feature on VLANs.

NOTE: Although you enable this feature on VLANs, the MAC move limitation pertains to the number of movements for each individual MAC address rather than the total number of MAC address moves in the VLAN. For example, If the MAC move limit is set to 1, the switch allows an unlimited number of MAC address movements within the VLAN as long as the same MAC address does not move more than once.

You configure MAC move limiting per VLAN, not per interface (port). In the default configuration, the number of MAC moves permitted is unlimited.

You can choose to have one of the following actions performed when the MAC move limit is exceeded:

- **drop**—Drop the packet and generate an alarm, an SNMP trap, or a system log entry. This is the default.
- **log**—Do not drop the packet but generate an alarm, an SNMP trap, or a system log entry.
- **none**—Take no action.
- **shutdown**—Disable the interfaces in the VLAN and generate an alarm. If you have configured the switch with the **port-error-disable** statement, the disabled interfaces recover automatically upon expiration of the specified disable timeout. If you have not configured the switch for autorecovery from port error disabled conditions, you can bring up the disabled interfaces by running the **clear ethernet-switching port-error** command.

To configure a MAC move limit for MAC addresses within a specific VLAN or for MAC addresses within all VLANs, using the CLI:

- On a single VLAN: To limit the number of MAC address movements that can be made by an individual MAC address within the VLAN **employee-vlan**, set a MAC move limit of **5**:

  ```
  [edit ethernet-switching-options secure-access-port]
  user@switch# set vlan employee–vlan mac-move-limit 5
  ```

  The action is not specified, so the switch performs the default action **drop** if it tracks that an individual MAC address within the **employee-vlan** has moved more than 5 times within one second.

- On all VLANs: To limit the number of MAC movements that can be made by individual MAC addresses within all VLANs, set a MAC move limit of **5**:

  ```
  [edit ethernet-switching-options secure-access-port]
  user@switch# set vlan all mac–move-limit 5
  ```

  The action is not specified, so the switch performs the default action **drop** if it tracks that an individual MAC address within any of the VLANs has moved more than 5 times within one second.

Related Documentation

- Configuring MAC Move Limiting (J-Web Procedure) on page 2641
- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Verifying That MAC Move Limiting Is Working Correctly on page 2661
- Monitoring Port Security on page 2653
- Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure) on page 2516
- Understanding MAC Limiting and MAC Move Limiting for Port Security on J-EX Series Switches on page 2557

Configuring MAC Move Limiting (J-Web Procedure)

MAC move limiting detects MAC address movement and MAC address spoofing on access ports. MAC address movements are tracked, and if a MAC address moves more than the configured number of times within one second, the configured (or default) action is performed. You enable this feature on VLANs.

NOTE: Although you enable this feature on VLANs, the MAC move limitation pertains to the number of movements for each individual MAC address rather than the total number of MAC address moves in the VLAN. For example, If the MAC move limit is set to 1, the switch allows an unlimited number of MAC address movements within the VLAN as long as the same MAC address does not move more than once.

In the default configuration, the MAC move limit within each VLAN is unlimited; the default action that the switch will take if the specified MAC move limit is exceeded is **drop**.

To enable MAC move limiting for MAC addresses within one or more VLANs by using the J-Web interface:

1. Select **Configure>Security>Port Security**.
2. Select one or more VLANs from the **VLAN List**.
3. Click the **Edit** button. If a message appears asking whether you want to enable port security, click **Yes**.
4. To set a MAC move limit:
    1. Type a limit value in the **MAC Movement** box.
    2. Select an action from the **MAC Movement Action** box (optional). The switch takes this action when an individual MAC address exceeds the MAC move limit. If you do not select an action, the switch applies the default action, **drop**.

       Select one:

       - Log—Generate a system log entry, an SNMP trap, or an alarm.
       - Drop—Drop the packets and generate a system log entry, an SNMP trap, or an alarm. (Default)
       - Shutdown—Shut down the VLAN and generate an alarm. You can mitigate the effect of this option by configuring the switch for autorecovery from the disabled state and specifying a **disable timeout** value. See “Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure)” on page 2516. If you have not configured autorecovery from the disabled state, you can bring up the interfaces by running the **clear ethernet-switching port-error** command.
       - None— No action to be taken.

3. Click **OK**.

5. Click **OK** after the configuration has been successfully delivered.

NOTE: You can enable or disable port security on the switch at any time by clicking the Activate or Deactivate button on the Port Security Configuration page. If security status is shown as Disabled when you try to edit settings for any VLANs, a message asking whether you want to enable port security appears.

Related Documentation

- Configuring MAC Move Limiting (CLI Procedure) on page 2639
- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Verifying That MAC Move Limiting Is Working Correctly on page 2661
- Monitoring Port Security on page 2653
- Understanding MAC Limiting and MAC Move Limiting for Port Security on J-EX Series Switches on page 2557

Setting the none Action on an Interface to Override a MAC Limit Applied to All Interfaces (CLI Procedure)

If you set a MAC limit in your port security settings to apply to all interfaces on the J-EX Series switch, you can override that setting for a particular interface by specifying action **none**.

To use the **none** action to override a MAC limit setting:

1. Set the MAC limit—for example, a limit of **5** with action **drop**:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface all mac-limit 5 action drop
```

2. Then change the action for one interface (here, **ge-0/0/2**) with this command. You don't need to specify a limit value.

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface ge–0/0/2 mac-limit action none
```

Related Documentation

- Configuring MAC Limiting (CLI Procedure) on page 2635
- Configuring MAC Limiting (J-Web Procedure) on page 2637
- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Verifying That MAC Limiting Is Working Correctly on page 2657

Configuring IP Source Guard (CLI Procedure)

You can use the IP source guard access port security feature on J-EX Series switches to mitigate the effects of source IP address spoofing and source MAC address spoofing. If IP source guard determines that a host connected to an access interface has sent a packet with an invalid source IP address or source MAC address in the packet header, it ensures that the switch does not forward the packet—that is, the packet is discarded.

You enable the IP source guard feature on VLANs. You can enable it on a specific VLAN, on all VLANs, or on a VLAN range.

NOTE: IP source guard applies only to access interfaces and only to untrusted interfaces. If you enable IP source guard on a VLAN that includes trunk interfaces or an interface set to **dhcp-trusted**, the CLI shows an error when you try to commit the configuration.

Before you configure IP source guard, be sure that you have:

Enabled DHCP snooping on the VLAN or VLANs on which you will configure IP source guard. See “Enabling DHCP Snooping (CLI Procedure)” on page 2630.

To enable IP source guard on a VLAN, all VLANs, or a VLAN range (a series of tagged VLANs) by using the CLI:

NOTE: Replace values displayed in italics with values for your configuration.

- On a specific VLAN:

```
[edit ethernet-switching-options secure-access port]
user@switch#set vlan default ip-source-guard
```

- On all VLANs:

```
[edit ethernet-switching-options secure-access port]
user@switch# set vlan all ip-source-guard
```

- On a VLAN range:

  1. Set the VLAN range (the VLAN name is **employee**):

```
[edit vlans]
user@switch# set employeevlan-range 100-101
```

  2. Associate an interface with a VLAN-range number (**100** in the following example) and set the port mode to **access**:

```
[edit interfaces]
user@switch# set ge-0/0/6 unit 0 family ethernet-switching port-mode access vlan
members100
```

  3. Enable IP source guard on the VLAN **employee**:

```
[edit ethernet-switching-options secure-access port]
user@switch# set vlan employee ip-source-guard
```

NOTE: You can use the **no-ip-source-guard** statement to disable IP source guard for a specific VLAN after you have enabled the feature for all VLANs.

To view results of the configuration steps before committing the configuration, type the **show** command at the user prompt.

To commit these changes to the active configuration, type the **commit** command at the user prompt.

Related Documentation

- Verifying That IP Source Guard Is Working Correctly on page 2662
- Example: Configuring IP Source Guard on a Data VLAN That Shares an Interface with a Voice VLAN on page 2608
- Example: Configuring IP Source Guard with Other J-EX Series Switch Features to Mitigate Address-Spoofing Attacks on Untrusted Access Interfaces on page 2600
- Understanding IP Source Guard for Port Security on J-EX Series Switches on page 2563

Configuring Static IP Addresses for DHCP Bindings on Access Ports (CLI Procedure)

You can add static (fixed) IP addresses and bind them to fixed MAC addresses in the DHCP snooping database. These bindings are labeled as “static” in the database, while those bindings that have been added through the process of DHCP snooping are labeled “dynamic.”

To configure a static IP address/MAC address binding in the DHCP snooping database (replace **ge-0/0/2**, **10.0.10.12**, **data-vlan**, and **00:05:85:3A:82:80** with values for your configuration):

```
[edit ethernet-switching-options secure-access-port]
user@switch# set interface ge-0/0/2 static-ip 10.0.10.12 vlan data-vlan mac 00:05:85:3A:82:80
```

To view results of the configuration steps before committing the configuration, type the **show** command at the user prompt.

To commit these changes to the active configuration, type the **commit** command at the user prompt.

Related Documentation

- Verifying That DHCP Snooping Is Working Correctly on page 2654
- Understanding DHCP Snooping for Port Security on J-EX Series Switches on page 2549

Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure)

You can use DHCP option 82, also known as the DHCP relay agent information option, to help protect the J-EX Series switch against attacks such as spoofing (forging) of IP addresses and MAC addresses, and DHCP IP address starvation. Option 82 provides information about the network location of a DHCP client, and the DHCP server uses this information to implement IP addresses or other parameters for the client.

You can configure the DHCP option 82 feature in two topologies:

- The switch functions as a relay agent when the DHCP clients or the DHCP server is connected to the switch through a Layer 3 interface. On the switch, these interfaces are configured as routed VLAN interfaces, or RVIs. The switch relays the clients' requests to the server and then forwards the server's replies to the clients. This topic describes this configuration.
- The switch, DHCP clients, and DHCP server are all on the same VLAN. The switch forwards the clients' requests to the server and forwards the server's replies to the clients. This configuration is described in “Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)” on page 2649.

Before you configure DHCP option 82 on the switch, perform these tasks:

- Connect and configure the DHCP server.

  NOTE: Your DHCP server must be configured to accept DHCP option 82. If the server is not configured for DHCP option 82, the server does not use the DHCP option 82 information in the requests sent to it when it formulates its reply messages.

- Configure the VLAN on the switch and associate the interfaces on which the clients connect to the switch with that VLAN.
- Configure the routed VLAN interface (RVI) to allow the switch to relay packets to the server and receive packets from the server. See “Configuring Routed VLAN Interfaces (CLI Procedure)” on page 1137.
- Configure the switch as a BOOTP relay agent. See “DHCP/BOOTP Relay for J-EX Series Switches Overview” on page 446.

To configure DHCP option 82:

NOTE: Replace values displayed in italics with values for your configuration.

1. Specify DHCP option 82 for the BOOTP server:

   - On all interfaces that connect to the server:

     ```
     [edit forwarding-options helpers bootp]
     user@switch# set dhcp-option82
     ```

   - On a specific interface that connects to the server:

     ```
     [edit forwarding-options helpers bootp]
     user@switch# set interface ge-0/0/10 dhcp-option82
     ```

     The remaining steps are optional. They show configurations for all interfaces; include the specific interface designation to configure any of the following options on a specific interface:

2. To configure a prefix for the circuit ID suboption (the prefix is always the hostname of the switch):

   ```
   [edit forwarding-options helpers bootp]
   user@switch# set dhcp-option82 circuit-id prefix hostname
   ```

3. To specify that the circuit ID suboption value contains the interface description rather than the interface name (the default):

   ```
   [edit forwarding-options helpers bootp]
   user@switch# set dhcp-option82 circuit-id use-interface-description
   ```

4. To specify that the circuit ID suboption value contains the VLAN ID rather than the VLAN name (the default):

   ```
   [edit forwarding-options helpers bootp]
   user@switch# set dhcp-option82 circuit-id use-vlan-id
   ```

5. To specify that the remote ID suboption is included in the DHCP option 82 information:

   ```
   [edit forwarding-options helpers bootp]
   user@switch# set dhcp-option82 remote-id
   ```

6. To configure a prefix for the remote ID suboption (here, the prefix is the MAC address of the switch):

   ```
   [edit forwarding-options helpers bootp]
   user@switch# set dhcp-option82 remote-id prefix mac
   ```

7. To specify that the prefix for the remote ID suboption is the hostname of the switch rather than the MAC address of the switch (the default):

   ```
   [edit forwarding-options helpers bootp]
   user@switch# set dhcp-option82 remote-id prefix hostname
   ```

8. To specify that the remote ID suboption value contains the interface description:

   ```
   [edit forwarding-options helpers bootp]
   user@switch# set dhcp-option82 remote-id use-interface-description
   ```

9. To specify that the remote ID suboption value contains a character string:

```
[edit forwarding-options helpers bootp]
user@switch# set dhcp-option82 remote-id use-string mystring
```

10. To configure a vendor ID suboption and use the default value (the default value is **Juniper**), do not type a character string after the **vendor-id** option keyword:

```
[edit forwarding-options helpers bootp]
user@switch# set dhcp-option82 vendor-id
```

11. To specify that the vendor ID suboption value contains a character string value that you specify rather than **Juniper** (the default):

```
[edit forwarding-options helpers bootp]
user@switch# set dhcp-option82 vendor-id mystring
```

To view results of the configuration steps before committing the configuration, type the **show** command at the user prompt.

To commit these changes to the active configuration, type the **commit** command at the user prompt.

Related Documentation

- Example: Setting Up DHCP Option 82 with a J-EX Series Switch as Relay Agent Between Clients and a DHCP Server on page 2615
- [edit forwarding-options] Configuration Statement Hierarchy on page 43
- Understanding DHCP Option 82 for Port Security on J-EX Series Switches on page 2560
- RFC 3046, *DHCP Relay Agent Information Option*, at http://tools.ietf.org/html/rfc3046.

Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure)

You can use DHCP option 82, also known as the DHCP relay agent information option, to help protect the J-EX Series switch against attacks such as spoofing (forging) of IP addresses and MAC addresses, and DHCP IP address starvation. Option 82 provides information about the network location of a DHCP client, and the DHCP server uses this information to implement IP addresses or other parameters for the client.

You can configure the DHCP option 82 feature in two topologies:

- The switch, DHCP clients, and DHCP server are all on the same VLAN. The switch forwards the clients' requests to the server and forwards the server's replies to the clients. This topic describes this configuration.
- The switch functions as a relay agent when the DHCP clients or the DHCP server is connected to the switch through a Layer 3 interface. On the switch, these interfaces are configured as routed VLAN interfaces, or RVIs. The switch relays the clients' requests to the server and then forwards the server's replies to the clients. This configuration is described in “Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure)” on page 2646.

Before you configure DHCP option 82 on the switch, perform these tasks:

- Connect and configure the DHCP server.

  

  NOTE: Your DHCP server must be configured to accept DHCP option 82. If the server is not configured for DHCP option 82, the server does not use the DHCP option 82 information in the requests sent to it when it formulates its reply messages.

- Configure a VLAN on the switch and associate the interfaces on which the clients and the server connect to the switch with that VLAN.

To configure DHCP option 82:

NOTE: Replace values displayed in italics with values for your configuration.

1. Specify DHCP option 82 for all VLANs associated with the switch or for a specified VLAN. (You can also configure the feature for a VLAN range.)

   - On a specific VLAN:

     ```
     [edit ethernet-switching-options secure-access-port]
     user@switch# set vlan employee dhcp-option82
     ```

   - On all VLANs:

     ```
     [edit ethernet-switching-options secure-access-port]
     user@switch# set vlan all dhcp-option82
     ```

   The remaining steps are optional.

2. To configure a prefix for the circuit ID suboption (the prefix is always the hostname of the switch):

   ```
   [edit ethernet-switching-options secure-access-port]
   user@switch# set vlan employee dhcp-option82 circuit-id prefix hostname
   ```

3. To specify that the circuit ID suboption value contains the interface description rather than the interface name (the default):

   ```
   [edit ethernet-switching-options secure-access-port]
   user@switch# set vlan employee dhcp-option82 circuit-id use-interface-description
   ```

4. To specify that the circuit ID suboption value contains the VLAN ID rather than the VLAN name (the default):

   ```
   [edit ethernet-switching-options secure-access-port]
   user@switch# set vlan employee dhcp-option82 circuit-id use-vlan-id
   ```

5. To specify that the remote ID suboption is included in the DHCP option 82 information:

   ```
   [edit ethernet-switching-options secure-access-port]
   user@switch# set vlan employee dhcp-option82 remote-id
   ```

6. To configure a prefix for the remote ID suboption (here, the prefix is the MAC address of the switch):

   ```
   [edit ethernet-switching-options secure-access-port]
   user@switch# set vlan employee dhcp-option82 remote-id prefix mac
   ```

7. To specify that the prefix for the remote ID suboption is the hostname of the switch rather than the MAC address of the switch (the default):

   ```
   [edit ethernet-switching-options secure-access-port]
   user@switch# set vlan employee dhcp-option82 remote-id prefix hostname
   ```

8. To specify that the remote ID suboption value contains the interface description:

   ```
   [edit ethernet-switching-options secure-access-port]
   user@switch# set vlan employee dhcp-option82 remote-id use-interface-description
   ```

9. To specify that the remote ID suboption value contains a character string:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee dhcp-option82 remote-id use-string mystring
```

10. To configure a vendor ID suboption and use the default value (the default value is **Juniper**), do not type a character string after the **vendor-id** option keyword:

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee dhcp-option82 vendor-id
```

11. To specify that the vendor ID suboption value contains a character string value that you specify rather than **Juniper** (the default):

```
[edit ethernet-switching-options secure-access-port]
user@switch# set vlan employee dhcp-option82 vendor-id mystring
```

To view results of the configuration steps before committing the configuration, type the **show** command at the user prompt.

To commit these changes to the active configuration, type the **commit** command at the user prompt.

Related Documentation

- Example: Setting Up DHCP Option 82 on a J-EX Series Switch with No Relay Agent Between Clients and DHCP Server on page 2617
- Understanding DHCP Option 82 for Port Security on J-EX Series Switches on page 2560
- RFC 3046, *DHCP Relay Agent Information Option*, at http://tools.ietf.org/html/rfc3046.

Configuring Proxy ARP (CLI Procedure)

You can configure proxy Address Resolution Protocol (ARP) on your J-EX Series switch to enable the switch to respond to ARP queries for network addresses by offering its own media access control (MAC) address. With proxy ARP enabled, the switch captures and routes traffic to the intended destination.

To configure proxy ARP on a single interface:

```
[edit interfaces]
user@switch# set ge-0/0/3 unit 0 proxy-arp restricted
```

BEST PRACTICE: We recommend that you configure proxy ARP in restricted mode. In restricted mode, the switch is not a proxy if the source and target IP addresses are on the same subnet. If you use unrestricted mode, disable gratuitous ARP requests on the interface to avoid the situation of the switch's response to a gratuitous ARP request appearing to the host to be an indication of an IP conflict:

To configure proxy ARP on a routed VLAN interface (RVI):

```
[edit interfaces]
user@switch# set vlan unit 100 proxy-arp restricted
```

Related Documentation

- Example: Configuring Proxy ARP on a J-EX Series Switch on page 2621
- Verifying That Proxy ARP Is Working Correctly on page 1164
- Configuring Routed VLAN Interfaces (CLI Procedure) on page 1137

Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure)

An Ethernet switching access interface on a J-EX Series switch might shut down or be disabled as a result of one of the following port-security or storm-control configurations:

- MAC limiting—**mac-limit** statement is configured with action **shutdown**.
- MAC move limiting—**mac-move-limit** statement is configured with action **shutdown**.
- Storm control—**storm-control** statement is configured with the action **shutdown**.

You can configure the switch to automatically restore the disabled interfaces to service after a specified period of time. Autorecovery applies to all the interfaces that have been disabled due to MAC limiting, MAC move limiting, or storm control errors.

NOTE: You must specify the disable timeout value for the interfaces to recover automatically. There is no default disable timeout. If you do not specify a timeout value, you need to use the clear ethernet-switching port-error command to clear the errors and restore the interfaces or the specified interface to service.

To configure autorecovery from the disabled state due to MAC limiting, MAC move limiting, or storm control shutdown actions:

```
[edit ethernet-switching-options]
user@switch# set port-error-disable disable-timeout 60
```

Related Documentation

- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Configuring MAC Limiting (CLI Procedure) on page 2635
- Example: Configuring Storm Control to Prevent Network Outages on J-EX Series Switches on page 2513
- Understanding MAC Limiting and MAC Move Limiting for Port Security on J-EX Series Switches on page 2557
- Understanding Storm Control on J-EX Series Switches on page 2511

CHAPTER 96

Verifying Port Security

Monitoring Port Security

Purpose Use the monitoring functionality to view these port security details:

- DHCP snooping database for a VLAN or all VLANs
- ARP inspection details for all interfaces

Action To monitor port security in the J-Web interface, select **Monitor > Security > Port Security**.

To monitor and manipulate the DHCP snooping database and ARP inspection statistics in the CLI, enter the following commands:

- **show dhcp snooping binding**
- **clear dhcp snooping binding**—In addition to clearing the whole database, you can clear database entries for specified VLANs or MAC addresses.
- **show arp inspection statistics**
- **clear arp inspection statistics**

Meaning The J-Web Port Security Monitoring page comprises two sections:

- DHCP Snooping—Displays the DHCP snooping database for all the VLANs for which DHCP snooping is enabled. To view the DHCP snooping database for a specific VLAN, select the specific VLAN from the list.

- ARP Inspection—Displays the ARP inspection details for all interfaces. The information includes details of the number of packets that passed ARP inspection and the number of packets that failed the inspection. The pie chart graphically represents these statistics when you select an interface. To view ARP inspection statistics for a specific interface, select the interface from the list.

You have the following options on the page:

- Clear ALL—Clears the DHCP snooping database, either for all VLANs if the option **ALL** has been selected in the Select VLANs list or for the specific VLAN that has been selected in that list.
- Clear—Deletes a specific IP address from the DHCP snooping database.

To clear ARP statistics on the page, click **Clear All** in the ARP Statistics section.

Use the CLI commands to show and clear DHCP snooping database and ARP inspection statistics details.

Related Documentation

- Configuring Port Security (CLI Procedure) on page 2626
- Configuring Port Security (J-Web Procedure) on page 2627
- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569

Verifying That DHCP Snooping Is Working Correctly

Purpose Verify that DHCP snooping is working on the switch and that the DHCP snooping database is correctly populated with both dynamic and static bindings.

Action Send some DHCP requests from network devices (here they are DHCP clients) connected to the switch.

Display the DHCP snooping information when the interface on which the DHCP server connects to the switch is trusted. The following output results when requests are sent from the MAC addresses and the server has provided the IP addresses and leases:

```
user@switch>  show dhcp snooping binding
DHCP Snooping Information:
MAC address        IP address  Lease (seconds) Type     VLAN      Interface

00:05:85:3A:82:77  192.0.2.17  600             dynamic  employee  ge-0/0/1.0
00:05:85:3A:82:79  192.0.2.18  653             dynamic  employee  ge-0/0/1.0
00:05:85:3A:82:80  192.0.2.19  720             dynamic  employee  ge-0/0/2.0
00:05:85:3A:82:81  192.0.2.20  932             dynamic  employee  ge-0/0/2.0
00:05:85:3A:82:83  192.0.2.21  1230            dynamic  employee  ge-0/0/2.0
00:05:85:27:32:88  192.0.2.22  –               static   data      ge-0/0/4.0
```

Meaning When the interface on which the DHCP server connects to the switch has been set to trusted, the output (see preceding sample) shows, for each MAC address, the assigned

IP address and lease time—that is, the time, in seconds, remaining before the lease expires. Static IP addresses have no assigned lease time. The statically configured entry never expires.

If the DHCP server had been configured as untrusted, no entries would be added to the DHCP snooping database and nothing would be shown in the output of the **show dhcp snooping binding** command.

Related Documentation

- Enabling DHCP Snooping (CLI Procedure) on page 2630
- Enabling DHCP Snooping (J-Web Procedure) on page 2631
- Configuring Static IP Addresses for DHCP Bindings on Access Ports (CLI Procedure) on page 2645
- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Example: Configuring DHCP Snooping, DAI , and MAC Limiting on a J-EX Series Switch with Access to a DHCP Server Through a Second Switch on page 2593
- Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks on page 2586
- Monitoring Port Security on page 2653
- Troubleshooting Port Security on page 2665

Verifying That a Trusted DHCP Server Is Working Correctly

Purpose Verify that a DHCP trusted server is working on the switch. See what happens when the DHCP server is trusted and then untrusted.

Action Send some DHCP requests from network devices (here they are DHCP clients) connected to the switch.

Display the DHCP snooping information when the interface on which the DHCP server connects to the switch is trusted. The following output results when requests are sent from the MAC addresses and the server has provided the IP addresses and leases:

```
user@switch>  show dhcp snooping binding
DHCP Snooping Information:
MAC Address         IP Address    Lease    Type     VLAN          Interface
-----------------   ----------    -----    ----     ----          ---------
00:05:85:3A:82:77    192.0.2.17    600    dynamic  employee-vlan   ge-0/0/1.0
00:05:85:3A:82:79    192.0.2.18    653    dynamic  employee-vlan   ge-0/0/1.0
00:05:85:3A:82:80    192.0.2.19    720    dynamic  employee-vlan   ge-0/0/2.0
00:05:85:3A:82:81    192.0.2.20    932    dynamic  employee-vlan   ge-0/0/2.0
00:05:85:3A:82:83    192.0.2.21   1230    dynamic  employee-vlan   ge-0/0/2.0
00:05:85:27:32:88    192.0.2.22   3200    dynamic  employee-vlan   ge-0/0/2.0
```

Meaning When the interface on which the DHCP server connects to the switch has been set to trusted, the output (see preceding sample) shows, for each MAC address, the assigned

IP address and lease time—that is, the time, in seconds, remaining before the lease expires.

If the DHCP server had been configured as untrusted, no entries would be added to the DHCP snooping database and nothing would be shown in the output of the **show dhcp snooping binding** command.

Related Documentation

- Enabling a Trusted DHCP Server (CLI Procedure) on page 2632
- Enabling a Trusted DHCP Server (J-Web Procedure) on page 2632
- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Example: Configuring a DHCP Server Interface as Untrusted to Protect the Switch from Rogue DHCP Server Attacks on page 2579
- Monitoring Port Security on page 2653
- Troubleshooting Port Security on page 2665

Verifying That DAI Is Working Correctly

Purpose Verify that dynamic ARP inspection (DAI) is working on the switch.

Action Send some ARP requests from network devices connected to the switch.

Display the DAI information:

```
user@switch> show arp inspection statistics
ARP inspection statistics:
Interface        Packets received  ARP inspection pass   ARP inspection failed
---------------  ---------------   -------------------- ---------------------
ge-0/0/1.0                     7                     5                     2
ge-0/0/2.0                    10                    10                     0
ge-0/0/3.0                    12                    12                     0
```

Meaning The sample output shows the number of ARP packets received and inspected per interface, with a listing of how many packets passed and how many failed the inspection on each interface. The switch compares the ARP requests and replies against the entries in the DHCP snooping database. If a MAC address or IP address in the ARP packet does not match a valid entry in the database, the packet is dropped.

Related Documentation

- Enabling Dynamic ARP Inspection (CLI Procedure) on page 2633
- Enabling Dynamic ARP Inspection (J-Web Procedure) on page 2634
- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Example: Configuring DHCP Snooping, DAI , and MAC Limiting on a J-EX Series Switch with Access to a DHCP Server Through a Second Switch on page 2593
- Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks on page 2586

- Monitoring Port Security on page 2653

Verifying That MAC Limiting Is Working Correctly

MAC limiting protects against flooding of the Ethernet switching table. MAC limiting sets a limit on the number of MAC addresses that can be learned on a single Layer 2 access interface (port).

The Junos OS provides two MAC limiting methods:

- Maximum number of dynamic MAC addresses allowed per interface—When the limit is exceeded, incoming packets with new MAC addresses are dropped.
- Specific "allowed" MAC addresses for the access interface—Any MAC address that is not in the list of configured addresses is not learned.

To verify MAC limiting configurations:

1. Verifying That MAC Limiting for Dynamic MAC Addresses Is Working Correctly on page 2657
2. Verifying That Allowed MAC Addresses Are Working Correctly on page 2658
3. Verifying Results of Various Action Settings When the MAC Limit Is Exceeded on page 2658
4. Customizing the Ethernet Switching Table Display to View Information for a Specific Interface on page 2660

Verifying That MAC Limiting for Dynamic MAC Addresses Is Working Correctly

Purpose Verify that MAC limiting for dynamic MAC addresses is working on the switch.

Action Display the MAC addresses that have been learned. The following sample output shows the results when two packets were sent from hosts on **ge-0/0/1** and five packets requests were sent from hosts on **ge-0/0/2**, with both interfaces set to a MAC limit of **4** with the action **drop**:

```
user@switch>  show ethernet-switching table
Ethernet-switching table:  7 entries, 6 learned
  VLAN                MAC address         Type         Age    Interfaces

  employee-vlan       *                   Flood          -    ge-0/0/2.0
  employee-vlan       00:05:85:3A:82:77   Learn          0    ge-0/0/1.0
  employee-vlan       00:05:85:3A:82:79   Learn          0    ge-0/0/1.0
  employee-vlan       00:05:85:3A:82:80   Learn          0    ge-0/0/2.0
  employee-vlan       00:05:85:3A:82:81   Learn          0    ge-0/0/2.0
  employee-vlan       00:05:85:3A:82:83   Learn          0    ge-0/0/2.0
  employee-vlan       00:05:85:3A:82:85   Learn          0    ge-0/0/2.0
```

Meaning The sample output shows that with a MAC limit of **4** for each interface, the packet for a fifth MAC address on **ge-0/0/2** was dropped because it exceeded the MAC limit. The address was not learned, and thus an asterisk (*) rather than an address appears in the **MAC address** column in the first line of the sample output.

Verifying That Allowed MAC Addresses Are Working Correctly

Purpose Verify that allowed MAC addresses are working on the switch.

Action Display the MAC cache information after allowed MAC addresses have been configured on an interface. The following sample shows the MAC cache after 5 allowed MAC addresses had been configured on interface **ge/0/0/2**. In this instance, the interface was also set to a dynamic MAC limit of **4** with action **drop**.

```
user@switch>  show ethernet-switching table
Ethernet-switching table:  5 entries, 4 learned
  VLAN                MAC address         Type         Age    Interfaces

  employee-vlan       00:05:85:3A:82:80   Learn          0    ge-0/0/2.0
  employee-vlan       00:05:85:3A:82:81   Learn          0    ge-0/0/2.0
  employee-vlan       00:05:85:3A:82:83   Learn          0    ge-0/0/2.0
  employee-vlan       00:05:85:3A:82:85   Learn          0    ge-0/0/2.0
  employee-vlan       *                   Flood          -    ge-0/0/2.0
```

Meaning Because the MAC limit value for this interface had been set to **4**, only four of the five configured allowed addresses were learned and thus added to the MAC cache. Because the fifth address was not learned, an asterisk (*) rather than an address appears in the **MAC address** column in the last line of the sample output.

Verifying Results of Various Action Settings When the MAC Limit Is Exceeded

Purpose Verify the results provided by the various action settings for MAC limits—**drop**, **log**, **none**, and **shutdown**—when the limits are exceeded.

Action Display the results of the various action settings.

NOTE: You can view log messages by using the **show log messages** command. You can also have the log messages displayed by configuring the monitor start messages with the **monitor start messages** command.

- **drop** action—For MAC limiting configured with a **drop** action and with the MAC limit set to **5**:

```
user@switch>  show ethernet-switching table

Ethernet-switching table: 6 entries, 5 learned
  VLAN                MAC address         Type         Age    Interfaces


  employee-vlan       *                   Flood          -    ge-0/0/2.0

  employee-vlan       00:05:85:3A:82:80   Learn          0    ge-0/0/2.0

  employee-vlan       00:05:85:3A:82:81   Learn          0    ge-0/0/2.0

  employee-vlan       00:05:85:3A:82:83   Learn          0    ge-0/0/2.0
```

```
  employee-vlan        00:05:85:3A:82:85   Learn          0    ge-0/0/2.0

  employee-vlan        00:05:85:3A:82:88   Learn          0    ge-0/0/2.0
```

- **log** action—For MAC limiting configured with a **log** action and with MAC limit set to **5**:

```
user@switch> show ethernet-switching table

Ethernet-switching table: 74 entries, 73 learned
  VLAN                 MAC address         Type         Age    Interfaces

  employee-vlan        *                   Flood          -    ge-0/0/2.0

  employee-vlan        00:05:85:3A:82:80   Learn          0    ge-0/0/2.0

  employee-vlan        00:05:85:3A:82:81   Learn          0    ge-0/0/2.0

  employee-vlan        00:05:85:3A:82:82   Learn          0    ge-0/0/2.0

  employee-vlan        00:05:85:3A:82:83   Learn          0    ge-0/0/2.0

  employee-vlan        00:05:85:3A:82:84   Learn          0    ge-0/0/2.0

  employee-vlan        00:05:85:3A:82:85   Learn          0    ge-0/0/2.0

  employee-vlan        00:05:85:3A:82:87   Learn          0    ge-0/0/2.0

  employee-vlan        00:05:85:3A:82:88   Learn          0    ge-0/0/2.0

  . . .
```

- **shutdown** action—For MAC limiting configured with a **shutdown** action and with MAC limit set to **3**:

```
user@switch> show ethernet-switching table

Ethernet-switching table: 4 entries, 3 learned
  VLAN                 MAC address         Type         Age    Interfaces

  employee-vlan        *                   Flood          -    ge-0/0/2.0

  employee-vlan        00:05:85:3A:82:82   Learn          0    ge-0/0/2.0

  employee-vlan        00:05:85:3A:82:84   Learn          0    ge-0/0/2.0

  employee-vlan        00:05:85:3A:82:87   Learn          0    ge-0/0/2.0
```

- **none** action—If you set a MAC limit to apply to all interfaces on the switch, you can override that setting for a particular interface by specifying this action for that interface. See “Setting the none Action on an Interface to Override a MAC Limit Applied to All Interfaces (CLI Procedure)” on page 2642.

Meaning For the **drop** action results—The sixth MAC address exceeded the MAC limit. The request packet for that address was dropped. Only five MAC addresses have been learned on **ge-0/0/2**.

For the **log** action results—The sixth MAC address exceeded the MAC limit. No MAC addresses were blocked.

For the **shutdown** action results—The fourth MAC address exceeded the MAC limit. Only three MAC addresses have been learned on **ge-0/0/2**. The interface **ge-0/0/1** is shut down.

For more information about interfaces that have been shut down, use the **show ethernet-switching interfaces** command.

```
user@switch> show ethernet-switching interfaces
Interface    State  VLAN members        Tag   Tagging  Blocking


bme0.32770   down   mgmt                      untagged unblocked

ge-1/0/0.0   down   v1                        untagged MAC limit exceeded

ge-1/0/1.0   up     v1                        untagged unblocked

ge-1/0/2.0   up     v1                        untagged unblocked

me0.0        up     mgmt                      untagged unblocked
```

NOTE: You can configure the switch to recover automatically from this type of error condition by specifying the port-error-disable statement with a disable timeout value. The switch automatically restores the disabled interface to service when the disable timeout expires. The port-error-disable configuration does not apply to pre-existing error conditions. It impacts only error conditions that are detected after port-error-disable has been enabled and committed. To clear a pre-existing error condition and restore the interface to service, use the clear ethernet-switching port-error command.

Customizing the Ethernet Switching Table Display to View Information for a Specific Interface

Purpose You can use the **show ethernet-switching table** command to view information for a specific interface.

Action For example, to display the MAC addresses that have been learned on **ge-0/0/2** interface, type:

```
user@switch> show ethernet-switching table interface ge-0/0/2.0
Ethernet-switching table: 1 unicast entries

  VLAN              MAC address       Type         Age Interfaces

  v1                *                 Flood          - All-members

  v1                00:00:06:00:00:00 Learn          0 ge-2/0/0.0
```

Meaning The MAC limit value for **ge-0/0/2** had been set to **1**, and the output shows that only one MAC address was learned and thus added to the MAC cache. An asterisk (*) rather than an address appears in the **MAC address** column in the first line of the sample output.

Related Documentation

- Configuring MAC Limiting (CLI Procedure) on page 2635
- Configuring MAC Limiting (J-Web Procedure) on page 2637
- Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure) on page 2516
- Example: Configuring Allowed MAC Addresses to Protect the Switch from DHCP Snooping Database Alteration Attacks on page 2590
- Example: Configuring MAC Limiting, Including Dynamic and Allowed MAC Addresses, to Protect the Switch from Ethernet Switching Table Overflow Attacks on page 2576
- Example: Configuring MAC Limiting to Protect the Switch from DHCP Starvation Attacks on page 2583
- Monitoring Port Security on page 2653

Verifying That MAC Move Limiting Is Working Correctly

Purpose Verify that MAC move limiting is working on the switch.

Action Display the MAC addresses in the Ethernet switching table when MAC move limiting has been configured for a VLAN. The following sample shows the results after two of the hosts on **ge-0/0/2** sent packets after the MAC addresses for those hosts had moved to other interfaces more than five times in 1 second. The VLAN, **employee-vlan**, was set to a MAC move limit of **5** with the action **drop**:

```
user@switch> show ethernet-switching table

Ethernet-switching table:  7 entries, 4 learned
  VLAN                MAC address         Type         Age    Interfaces

  employee-vlan       00:05:85:3A:82:77   Learn          0    ge-0/0/1.0
  employee-vlan       00:05:85:3A:82:79   Learn          0    ge-0/0/1.0
  employee-vlan       00:05:85:3A:82:80   Learn          0    ge-0/0/2.0
  employee-vlan       00:05:85:3A:82:81   Learn          0    ge-0/0/2.0
  employee-vlan       *                   Flood          -    ge-0/0/2.0
  employee-vlan       *                   Flood          -    ge-0/0/2.0
```

Meaning The last two lines of the sample output show that MAC addresses for two hosts on **ge-0/0/2** were not learned, because the hosts had been moved back and forth from the original interfaces more than five times in 1 second.

NOTE: For descriptions of the results of the various action settings—**drop**, **log**, **none**, and **shutdown**—see "Verifying That MAC Limiting Is Working Correctly" on page 2657.

Related Documentation

- Configuring MAC Move Limiting (CLI Procedure) on page 2639
- Configuring MAC Move Limiting (J-Web Procedure) on page 2641
- Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure) on page 2516
- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Monitoring Port Security on page 2653

Verifying That IP Source Guard Is Working Correctly

Purpose Verify that IP source guard is enabled and is mitigating the effects of any source IP spoofing attacks on the J-EX Series switch.

Action Display the IP source guard database.

```
user@switch> show ip-source-guard
IP source guard information:
Interface    Tag  IP Address   MAC Address        VLAN

ge-0/0/12.0  0    10.10.10.7   00:30:48:92:A5:9D  vlan100

ge-0/0/13.0  0    10.10.10.9   00:30:48:8D:01:3D  vlan100

ge–0/0/13.0  100  *            *                  voice
```

Meaning The IP source guard database table contains the VLANs enabled for IP source guard, the untrusted access interfaces on those VLANs, the VLAN 802.1Q tag IDs if there are any, and the IP addresses and MAC addresses that are bound to one another. If a switch interface is associated with multiple VLANs and some of those VLANs are enabled for IP source guard and others are not, the VLANs that are not enabled for IP source guard have a star (*) in the **IP Address** and **MAC Address** fields. See the entry for the **voice** VLAN in the preceding sample output.

Related Documentation

- Configuring IP Source Guard (CLI Procedure) on page 2643

Verifying That Proxy ARP Is Working Correctly

Purpose Verify that the switch is sending proxy ARP messages.

Action List the system statistics for ARP:

```
user@switch> show system statistics arp
arp:
        198319 datagrams received
        45 ARP requests received
        12 ARP replies received
        2 resolution requests received
        2 unrestricted proxy requests
        0 restricted proxy requests
        0 received proxy requests
```

```
0 proxy requests not proxied
0 restricted-proxy requests not proxied
0 with bogus interface
0 with incorrect length
0 for non-IP protocol
0 with unsupported op code
0 with bad protocol address length
0 with bad hardware address length
0 with multicast source address
0 with multicast target address
0 with my own hardware address
168705 for an address not on the interface
0 with a broadcast source address
0 with source address duplicate to mine
29555 which were not for me
0 packets discarded waiting for resolution
4 packets sent after waiting for resolution
27 ARP requests sent
47 ARP replies sent
0 requests for memory denied
0 requests dropped on entry
0 requests dropped during retry
0 requests dropped due to interface deletion
0 requests on unnumbered interfaces
0 new requests on unnumbered interfaces
0 replies for from unnumbered interfaces
0 requests on unnumbered interface with non-subnetted donor
0 replies from unnumbered interface with non-subnetted donor
```

Meaning The statistics show that two proxy ARP requests were received, and the **proxy requests not proxied** field indicates that all the unproxied ARP requests received have been proxied by the switch.

Related Documentation

- Configuring Proxy ARP (CLI Procedure) on page 1153

Verifying That the Port Error Disable Setting Is Working Correctly

Purpose Verify that the port error disable setting is working as expected on MAC limited, MAC move limited and rate-limited interfaces on a J-EX Series switch.

Action Display information about interfaces:

```
user@switch> show ethernet-switching interfaces
Interface   State    VLAN members           Blocking
ge-0/0/0.0  up       T1122                  unblocked
ge-0/0/1.0  down     default                MAC limit exceeded
ge-0/0/2.0  down     default                MAC move limit exceeded
ge-0/0/3.0  down     default                Storm control in effect
ge-0/0/4.0  down     default                unblocked
ge-0/0/5.0  down     default                unblocked
ge-0/0/6.0  down     default                unblocked
ge-0/0/7.0  down     default                unblocked
ge-0/0/8.0  down     default                unblocked
ge-0/0/9.0  up       T111                   unblocked
ge-0/0/10.0 down     default                unblocked
ge-0/0/11.0 down     default                unblocked
ge-0/0/12.0 down     default                unblocked
```

```
ge-0/0/13.0 down     default              unblocked
ge-0/0/14.0 down     default              unblocked
ge-0/0/15.0 down     default              unblocked
ge-0/0/16.0 down     default              unblocked
ge-0/0/17.0 down     default              unblocked
ge-0/0/18.0 down     default              unblocked
ge-0/0/19.0 up       T111                 unblocked
ge-0/1/0.0  down     default              unblocked
ge-0/1/1.0  down     default              unblocked
ge-0/1/2.0  down     default              unblocked
ge-0/1/3.0  down     default              unblocked
```

Meaning The sample output from the **show ethernet-switching interfaces** command shows that three of the down interfaces specify the reason that the interface is disabled:

- **MAC limit exceeded**—The interface is temporarily disabled due to a **mac-limit** error. The disabled interface is automatically restored to service when the **disable-timeout** expires.
- **MAC move limit exceeded**—The interface is temporarily disabled due to a **mac-move-limit** error. The disabled interface is automatically restored to service when the **disable-timeout** expires.
- **Storm control in efffect** —The interface is temporarily disabled due to a **storm-control** error. The disabled interface is automatically restored to service when the **disable-timeout** expires.

Related Documentation

- Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure) on page 2516

CHAPTER 97

Troubleshooting Port Security

Troubleshooting Port Security

Troubleshooting issues for port security on J-EX Series switches:

MAC Addresses That Exceed the MAC Limit or MAC Move Limit Are Not Listed in the Ethernet Switching Table

Problem You see log messages telling you that the MAC limit or MAC move limit has been exceeded, but the specific offending MAC addresses that have been exceeding the limit are not listed in the Ethernet switching table.

Solution

1. Set the MAC limit or MAC move limit action to **log**.

```
[edit ethernet-switching-options secure-access port]
user@switch# set interface ge-0/0/2 mac-limit 5 action log
```

2. Allow some MAC address requests to come in.
3. View the entries in the Ethernet switching table:

```
user@switch> show ethernet-switching table
```

Multiple DHCP Server Packets Have Been Received on Untrusted Interfaces

Problem You see log messages that DHCP server packets were received on an untrusted interface—for example:

```
5 untrusted DHCPOFFER received, interface ge-0/0/0.0[65], vlan v1[10] server
ip/mac 12.12.12.1/00:00:00:00:01:12 offer ip/client mac
12.12.12.253/00:AA:BB:CC:DD:01
```

These messages can signal the presence of a malicious DHCP server on the network.

Solution Configure a firewall filter to block the IP address or MAC address of the malicious DHCP server. See “Configuring Firewall Filters (CLI Procedure)” on page 2779.

Related Documentation

- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Verifying That a Trusted DHCP Server Is Working Correctly on page 2655
- Verifying That MAC Limiting Is Working Correctly on page 2657
- Enabling a Trusted DHCP Server (CLI Procedure) on page 2632
- Configuring MAC Limiting (CLI Procedure) on page 2635

CHAPTER 98

Configuration Statements for Port Security

[edit ethernet-switching-options] Configuration Statement Hierarchy

```
ethernet-switching-options {
  analyzer {
    name {
      loss-priority priority;
      ratio number;
      input {
        ingress {
          interface (all | interface-name);
          vlan (vlan-id | vlan-name);
        }
        egress {
          interface (all | interface-name);
        }
      output {
        interface interface-name;
        vlan (vlan-id | vlan-name);
      }
    }
  }
  bpdu-block {
    disable-timeout timeout;
    interface (all | [interface-name]);
  }
  dot1q-tunneling {
    ether-type (0x8100 | 0x88a8 | 0x9100);
  }
  interfaces interface-name {
    no-mac-learning;
  }
  mac-notification {
    notification-interval seconds;
  }
  mac-table-aging-time seconds;
```

```
port-error-disable {
  disable-timeout timeout;
}
redundant-trunk-group {
  group-name name {
    interface interface-name <primary>;
  }
}
secure-access-port {
  dhcp-snooping-file {
    location local_pathname | remote_URL;
    timeout seconds;
    write-interval seconds;
  }
  interface (all | interface-name) {
    allowed-mac {
      mac-address-list;
    }
    (dhcp-trusted | no-dhcp-trusted );
    mac-limit limit action action;
    no-allowed-mac-log;
    static-ip  ip-address {
      vlan  vlan-name;
      mac mac-address;
    }
  }
  vlan (all | vlan-name) {
    (arp-inspection | no-arp-inspection );
    dhcp-option82 {
      circuit-id {
        prefix hostname;
        use-interface-description;
        use-vlan-id;
      }
      remote-id {
        prefix hostname | mac | none;
        use-interface-description;
        use-string string;
      }
      vendor-id [string];
    }
    (examine-dhcp | no-examine-dhcp );
    (ip-source-guard | no-ip-source-guard);
    mac-move-limit limit action action;
  }
}
storm-control {
  action-shutdown;
  interface (all | interface-name) {
    bandwidth bandwidth;
    no-broadcast;
    no-unknown-unicast;
  }
}
traceoptions {
```

```
    file filename <files number> <no-stamp> <replace> <size size> <world-readable |
      no-world-readable>;
    flag flag <disable>;
  }
  unknown-unicast-forwarding {
    vlan (all | vlan-name) {
      interface interface-name;
    }
  }
  voip {
    interface (all | [interface-name | access-ports]) {
      vlan vlan-name ;
      forwarding-class <assured-forwarding | best-effort | expedited-forwarding |
        network-control>;
    }
  }
}
```

Related Documentation

- Understanding Port Mirroring on J-EX Series Switches on page 3245
- Port Security for J-EX Series Switches Overview on page 2545
- Understanding BPDU Protection for STP, RSTP, and MSTP on J-EX Series Switches on page 1278
- Understanding Redundant Trunk Links on J-EX Series Switches on page 1049
- Understanding Storm Control on J-EX Series Switches on page 2511
- Understanding 802.1X and VoIP on J-EX Series Switches on page 2263
- Understanding Q-in-Q Tunneling on J-EX Series Switches on page 1051
- Understanding Unknown Unicast Forwarding on J-EX Series Switches on page 2512
- Understanding MAC Notification on J-EX Series Switches on page 1060

[edit forwarding-options] Configuration Statement Hierarchy

```
helpers {
  bootp {
    dhcp-option82 {
      circuit-id {
        prefix hostname;
        use-interface-description;
        use-vlan-id;
      }
      remote-id {
        prefix hostname | mac | none;
        use-interface-description;
        use-string string;
      }
      vendor-id <string>;
    }
    interface interface-name {
      dhcp-option82 {
        circuit-id {
```

```
            prefix hostname;
            use-interface-description;
            use-vlan-id;
          }
          remote-id {
            prefix hostname | mac | none;
            use-interface-description;
            use-string string;
          }
          vendor-id <string>;
        }
        source-address-giaddr;
      }
      source-address-giaddr;
    }
  }
```

Related Documentation

- Example: Setting Up DHCP Option 82 with a J-EX Series Switch as Relay Agent Between Clients and a DHCP Server on page 2615
- Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 2646
- Understanding DHCP Option 82 for Port Security on J-EX Series Switches on page 2560
- DHCP/BOOTP Relay for J-EX Series Switches Overview on page 446
- For more information about the **[edit forwarding-options]** hierarchy and all its options, see the *Junos OS Policy Framework Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.

allowed-mac

Syntax

```
allowed-mac {
  mac-address-list;
}
```

Hierarchy Level [edit **ethernet-switching-options secure-access-port interface** (all | *interface-name*)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify particular MAC addresses to be added to the MAC address cache.

NOTE: Although this configuration restricts the addresses that can be added to the MAC address cache, it does not block the switch from receiving Layer 2 control packets—such as Link Layer Discovery Protocol (LLDP) packets—transmitted from MAC addresses that are not specified in the list of allowed MAC addresses. Control packets do not undergo the MAC address check and they are therefore included in the statistics of packets received. However, they are not forwarded to another destination. They are trapped within the switch.

Default Allowed MAC addresses take precedence over dynamic MAC values that have been applied with the **mac-limit** statement.

Options *mac-address-list*—One or more MAC addresses configured as allowed MAC addresses for a specified interface or all interfaces.

Required Privilege Level routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation

- **mac-limit on page 2684**
- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Example: Configuring Allowed MAC Addresses to Protect the Switch from DHCP Snooping Database Alteration Attacks on page 2590
- Example: Configuring MAC Limiting, Including Dynamic and Allowed MAC Addresses, to Protect the Switch from Ethernet Switching Table Overflow Attacks on page 2576
- Example: Configuring MAC Limiting to Protect the Switch from DHCP Starvation Attacks on page 2583
- Configuring MAC Limiting (CLI Procedure) on page 2635
- Configuring MAC Limiting (J-Web Procedure) on page 2637

arp-inspection

Syntax (arp-inspection | no-arp-inspection);

Hierarchy Level [edit **ethernet-switching-options secure-access-port vlan** (all | *vlan-name*)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Perform dynamic ARP inspection (DAI) on all VLANs or on the specified VLAN.

- **arp-inspection**—Enable DAI.

  When ARP inspection is enabled, the switch logs invalid ARP packets that it receives on each interface, along with the sender's IP and MAC addresses.

- **no-arp-inspection**—Disable DAI.

Default Disabled.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Example: Configuring DHCP Snooping, DAI , and MAC Limiting on a J-EX Series Switch with Access to a DHCP Server Through a Second Switch on page 2593
- Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks on page 2586
- Enabling Dynamic ARP Inspection (CLI Procedure) on page 2633
- Enabling Dynamic ARP Inspection (J-Web Procedure) on page 2634

circuit-id

Syntax

```
circuit-id {
  prefix hostname;
  use-interface-description;
  use-vlan-id;
}
```

Hierarchy Level

[edit **ethernet-switching-options secure-access-port vlan** (all | *vlan-name*) **dhcp-option82**]
[edit **forwarding-options helpers bootp dhcp-option82**]
[edit **forwarding-options helpers bootp interface** *interface-name* **dhcp-option82**]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure the **circuit-id** suboption (suboption 1) of DHCP option 82 (the DHCP relay agent information option) in DHCP packets destined for a DHCP server. This suboption identifies the circuit (interface and/or VLAN) on which the DHCP request arrived.

The format of the **circuit-id** information for Gigabit Ethernet interfaces that use VLANs is ***interface-name*:*vlan-name*** . On a Layer 3 interface, the format is just ***interface-name*.**

The remaining statements are explained separately.

Default

If DCHP option 82 is enabled on the switch, the circuit ID is supplied by default in the format ***interface-name*:*vlan-name*** or, on a Layer 3 interface, just ***interface-name*** .

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Setting Up DHCP Option 82 on a J-EX Series Switch with No Relay Agent Between Clients and DHCP Server on page 2617
- Example: Setting Up DHCP Option 82 with a J-EX Series Switch as Relay Agent Between Clients and a DHCP Server on page 2615
- Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 2649
- Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 2646
- [edit forwarding-options] Configuration Statement Hierarchy on page 43
- RFC 3046, *DHCP Relay Agent Information Option*, at http://tools.ietf.org/html/rfc3046.

dhcp-option82

Syntax

```
dhcp-option82 {
  circuit-id {
    prefix hostname;
    use-interface-description;
    use-vlan-id;
  }
  remote-id {
    prefix hostname | mac | none;
    use-interface-description;
    use-string string;
  }
  vendor-id <string>;
}
```

Hierarchy Level

[edit **ethernet-switching-options secure-access-port vlan** (all | *vlan-name*)]
[edit **forwarding-options helpers bootp**]
[edit **forwarding-options helpers bootp interface** *interface-name*]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

When the switch receives a DHCP request from a DHCP client connected on one of the switch's interfaces, have the switch insert DHCP option 82 (also known as the DHCP relay agent information option) information in the DHCP request packet header before it forwards or relays the request to a DHCP server. The server uses the option 82 information, which provides details about the circuit and host the request came from, in formulating the reply; the server does not, however, make any changes to the option 82 information in the packet header. The switch receives the reply and then removes the DHCP option 82 information before forwarding the reply to the client.

The remaining statements are explained separately.

Default

Insertion of DHCP option 82 information is not enabled.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Setting Up DHCP Option 82 on a J-EX Series Switch with No Relay Agent Between Clients and DHCP Server on page 2617
- Example: Setting Up DHCP Option 82 with a J-EX Series Switch as Relay Agent Between Clients and a DHCP Server on page 2615
- Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 2649
- Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 2646
- [edit forwarding-options] Configuration Statement Hierarchy on page 43
- RFC 3046, *DHCP Relay Agent Information Option*, at http://tools.ietf.org/html/rfc3046.

dhcp-snooping-file

Syntax

```
dhcp-snooping-file {
  location local_pathname | remote_URL;
  timeout seconds;
  write-interval seconds;
}
```

Hierarchy Level [edit ethernet-switching-options secure-access-port]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify a local pathname or remote URL for the DHCP snooping database file to maintain persistence of IP-MAC bindings.

The remaining statements are explained separately.

Default The IP-MAC bindings in the DHCP snooping database file are not persistent. If the switch is rebooted, the bindings are lost.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Understanding DHCP Snooping for Port Security on J-EX Series Switches on page 2549

dhcp-trusted

Syntax (dhcp-trusted | no-dhcp-trusted);

Hierarchy Level [edit **ethernet-switching-options secure-access-port interface** (all | *interface-name*)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Allow DHCP responses from the specified interfaces (ports) or all interfaces.

- **dhcp-trusted**—Allow DHCP responses.
- **no-dhcp-trusted**—Deny DHCP responses.

Default Trusted for trunk ports, untrusted for access ports.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Example: Configuring a DHCP Server Interface as Untrusted to Protect the Switch from Rogue DHCP Server Attacks on page 2579
- Enabling a Trusted DHCP Server (CLI Procedure) on page 2632
- Enabling a Trusted DHCP Server (J-Web Procedure) on page 2632

disable-timeout

Syntax disable-timeout *timeout*;

Hierarchy Level [edit **ethernet-switching-options port-error-disable**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify how long the Ethernet switching interfaces remain in a disabled state due to MAC limiting, MAC move limiting, or storm control errors.

NOTE: If you modify the timeout value of an existing disable timeout, the new timeout value does not impact the timing of restoration to service of currently disabled interfaces that have been configured for automatic recovery. The new timeout value is applied only during the next occurrence of a port error.

You can bring up the currently disabled interfaces by running the clear ethernet-switching port-error command.

Default The disable timeout is not enabled.

Options ***timeout***—Time, in seconds, that the disabled state remains in effect. The disabled interface is automatically restored to service when the specified timeout value is reached.

Range: 10 through 3600 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Storm Control to Prevent Network Outages on J-EX Series Switches on page 2513
- Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure) on page 2516

ethernet-switching-options

Syntax

```
ethernet-switching-options {
  analyzer {
    name {
      loss-priority priority;
      ratio number;
      input {
        ingress {
          interface (all | interface-name);
          vlan (vlan-id | vlan-name);
        }
        egress {
          interface (all | interface-name);
        }
      output {
        interface interface-name;
        vlan (vlan-id | vlan-name);
      }
    }
  }
  bpdu-block {
    disable-timeout timeout;
    interface (all | [interface-name]);
  }
  dot1q-tunneling {
    ether-type (0x8100 | 0x88a8 | 0x9100);
  }
  interfaces interface-name {
    no-mac-learning;
  }
  mac-notification {
    notification-interval seconds;
  }
  mac-table-aging-time seconds;
  port-error-disable {
    disable-timeout timeout;
  }
  redundant-trunk-group {
    group-name name {
      interface interface-name <primary>;
      interface interface-name;
    }
  }
  secure-access-port {
    dhcp-snooping-file {
      location local_pathname | remote_URL;
      timeout seconds;
      write-interval seconds;
    }
    interface (all | interface-name) {
      allowed-mac {
        mac-address-list;
      }
```

```
        (dhcp-trusted | no-dhcp-trusted);
        mac-limit limit action action;
        no-allowed-mac-log;
        static-ip ip-address {
          vlan  vlan-name;
          mac mac-address;
        }
      }
      vlan (all | vlan-name) {
        (arp-inspection | no-arp-inspection);
        dhcp-option82 {
          circuit-id {
            prefix hostname;
            use-interface-description;
            use-vlan-id;
          }
          remote-id {
            prefix hostname | mac | none;
            use-interface-description;
            use-string string;
          }
          vendor-id [string];
        }
        (examine-dhcp | no-examine-dhcp);
        (ip-source-guard | no-ip-source-guard);
        mac-move-limit limit action action;
      }
    }
    storm-control {
      action-shutdown;
      interface (all | interface-name) {
        bandwidth bandwidth;
        no-broadcast;
        no-unknown-unicast;
      }
    }
    traceoptions {
      file filename <files number> <no-stamp> <replace> <size size> <world-readable |
        no-world-readable>;
      flag flag <disable>;
    }
    unknown-unicast-forwarding {
      vlan (all | vlan-name) {
        interface interface-name;
      }
    }
    voip {
      interface (all | [interface-name | access-ports]) {
        vlan vlan-name ;
        forwarding-class <assured-forwarding | best-effort | expedited-forwarding |
          network-control>;
      }
    }
  }
```

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure Ethernet switching options.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation

- Understanding Port Mirroring on J-EX Series Switches on page 3245
- Port Security for J-EX Series Switches Overview on page 2545
- Understanding BPDU Protection for STP, RSTP, and MSTP on J-EX Series Switches on page 1278
- Understanding Redundant Trunk Links on J-EX Series Switches on page 1049
- Understanding Storm Control on J-EX Series Switches on page 2511
- Understanding 802.1X and VoIP on J-EX Series Switches on page 2263
- Understanding Q-in-Q Tunneling on J-EX Series Switches on page 1051
- Understanding Unknown Unicast Forwarding on J-EX Series Switches on page 2512
- Understanding MAC Notification on J-EX Series Switches on page 1060

examine-dhcp

Syntax (examine-dhcp | no-examine-dhcp);

Hierarchy Level [edit **ethernet-switching-options secure-access-port vlan** (all | *vlan-name*)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable DHCP snooping on all VLANs or on the specified VLAN.

- **examine-dhcp**—Enable DHCP snooping.
- **no-examine-dhcp**—Disable DHCP snooping.

When DHCP snooping is enabled, the switch logs DHCP packets (DHCPOFFER, DHCPDECLINE, DHCPACK, and DHCPNAK packets) that it receives on untrusted ports. You can monitor the log for these messages, which can signal the presence of a malicious DHCP server on the network.

TIP: For Private VLANs (PVLANs), enable DHCP snooping on the primary VLAN. If you enable DHCP snooping only on a community VLAN, DHCP messages coming from PVLAN trunk ports are not snooped.

Default Disabled.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Example: Configuring DHCP Snooping, DAI , and MAC Limiting on a J-EX Series Switch with Access to a DHCP Server Through a Second Switch on page 2593
- Example: Configuring DHCP Snooping and DAI to Protect the Switch from ARP Spoofing Attacks on page 2586
- Enabling DHCP Snooping (CLI Procedure) on page 2630
- Enabling DHCP Snooping (J-Web Procedure) on page 2631

interface

Syntax

```
interface (all | interface-name) {
  allowed-mac {
    mac-address-list;
  }
  (dhcp-trusted | no-dhcp-trusted);
  mac-limit limit action action;
  no-allowed-mac-log;
  static-ip ip-address {
    vlan vlan-name;
    mac mac-address;
  }
}
```

Hierarchy Level [edit ethernet-switching-options secure-access-port]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply port security features to all interfaces or to the specified interface.

The statements are explained separately.

Options **all**—Apply port security features to all interfaces.

interface-name —Apply port security features to the specified interface.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Example: Configuring Allowed MAC Addresses to Protect the Switch from DHCP Snooping Database Alteration Attacks on page 2590
- Example: Configuring MAC Limiting, Including Dynamic and Allowed MAC Addresses, to Protect the Switch from Ethernet Switching Table Overflow Attacks on page 2576
- Example: Configuring MAC Limiting to Protect the Switch from DHCP Starvation Attacks on page 2583
- Example: Configuring a DHCP Server Interface as Untrusted to Protect the Switch from Rogue DHCP Server Attacks on page 2579
- Configuring MAC Limiting (CLI Procedure) on page 2635
- Enabling a Trusted DHCP Server (CLI Procedure) on page 2632
- Configuring Static IP Addresses for DHCP Bindings on Access Ports (CLI Procedure) on page 2645

ip-source-guard

Syntax (ip-source-guard | no-ip-source-guard);

Hierarchy Level [edit **ethernet-switching-options secure-access-port vlan** (all | *vlan-name*)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Perform IP source guard checking on packets sent from access interfaces. Validate source IP addresses and source MAC addresses on all VLANs or on the specified VLAN or VLAN range. Forward packets with valid addresses and drop those with invalid addresses.

- **ip-source-guard**—Enable IP source guard checking.
- **no-ip-source-guard**—Disable IP source guard checking.

Default Disabled.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring IP Source Guard on a Data VLAN That Shares an Interface with a Voice VLAN on page 2608
- Example: Configuring IP Source Guard with Other J-EX Series Switch Features to Mitigate Address-Spoofing Attacks on Untrusted Access Interfaces on page 2600
- Configuring IP Source Guard (CLI Procedure) on page 2643

mac

Syntax mac *mac-address*;

Hierarchy Level [edit **ethernet-switching-options secure-access-port interface** (all | *interface-name*) **static-ip** *ip-address* **vlan** *vlan-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Media access control (MAC) address, or hardware address, for the device connected to the specified interface.

Options ***mac-address*** —Value (in hexadecimal format) for address assigned to this device.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Static IP Addresses for DHCP Bindings on Access Ports (CLI Procedure) on page 2645

mac-limit

Syntax mac-limit *limit* action *action*;

Hierarchy Level [edit **ethernet-switching-options secure-access-port interface** (all | *interface-name*)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the number of MAC addresses to dynamically add to the MAC address cache for this access interface (port) and the action to be taken by the switch if the MAC address learning limit is exceeded on the interface (port).

Default The default action is **drop**.

Options ***limit***—Maximum number of MAC addresses.

action *action*—(Optional) Action to take when the MAC address limit is exceeded:

- **drop**—Drop the packet and generate an alarm, an SNMP trap, or a system log entry. This is the default.
- **log**—Do not drop the packet but generate an alarm, an SNMP trap, or a system log entry.
- **none**—No action.
- **shutdown**—Disable the interface and generate an alarm. If you have configured the switch with the **port-error-disable** statement, the disabled interface recovers automatically upon expiration of the specified disable timeout. If you have not configured the switch for autorecovery from port error disabled conditions, you can bring up the disabled interfaces by running the **clear ethernet-switching port-error** command.

Required Privilege Level routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation

- **allowed-mac on page 2671**
- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Example: Configuring MAC Limiting, Including Dynamic and Allowed MAC Addresses, to Protect the Switch from Ethernet Switching Table Overflow Attacks on page 2576
- Example: Configuring MAC Limiting to Protect the Switch from DHCP Starvation Attacks on page 2583
- Configuring MAC Limiting (CLI Procedure) on page 2635
- Configuring MAC Limiting (J-Web Procedure) on page 2637
- Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure) on page 2516

mac-move-limit

Syntax mac-move-limit *limit* action *action*;

Hierarchy Level [edit **ethernet-switching-options secure-access-port vlan** (all | *vlan-name*)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the number of times a MAC address can move to a new interface (port) in 1 second and the action to be taken by the switch if the MAC address move limit is exceeded.

Default The default move limit is unlimited. The default action is **drop**.

Options ***limit***—Maximum number of moves to a new interface per second.

action *action*—(Optional) Action to take when the MAC address move limit is reached:

- **drop**—Drop the packet and generate an alarm, an SNMP trap, or a system log entry. This is the default.
- **log**—Do not drop the packet but generate an alarm, an SNMP trap, or a system log entry.
- **none**—No action.
- **shutdown**—Disable the VLAN and generate an alarm. If you have configured the switch with the **port-error-disable** statement, the disabled interfaces recover automatically upon expiration of the specified disable timeout. If you have not configured the switch for autorecovery from port error disabled conditions, you can bring up the disabled interfaces by running the **clear ethernet-switching port-error** command.

Required Privilege Level routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation

- **mac-limit on page 2684**
- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Configuring MAC Move Limiting (CLI Procedure) on page 2639
- Configuring MAC Move Limiting (J-Web Procedure) on page 2641
- Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure) on page 2516

no-allowed-mac-log

Syntax no-allowed-mac-log;

Hierarchy Level [edit **ethernet-switching-options secure-access-port interface** (all | *interface-name*)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify that the switch does not log messages when it receives packets from invalid MAC addresses on an interface that has been configured for particular (allowed) MAC addresses.

Default The switch logs messages when it receives packets from invalid MAC addresses on an interface that has been configured for particular (allowed) MAC addresses.

Required Privilege Level routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation

- **allowed-mac on page 2671**
- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Example: Configuring Allowed MAC Addresses to Protect the Switch from DHCP Snooping Database Alteration Attacks on page 2590
- Example: Configuring MAC Limiting to Protect the Switch from DHCP Starvation Attacks on page 2583
- Configuring MAC Limiting (CLI Procedure) on page 2635
- Configuring MAC Limiting (J-Web Procedure) on page 2637

no-gratuitous-arp-request

Syntax no-gratuitous-arp-request;

Hierarchy Level [edit **interfaces** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the switch not to respond to gratuitous ARP requests. You can disable responses to gratuitous ARP requests on both Layer 2 Ethernet switching interfaces and routed VLAN interfaces (RVIs).

Default Gratuitous ARP responses are enabled on all Ethernet switching interfaces and RVIs.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Unrestricted Proxy ARP on a J-EX Series Switch on page 2621
- Configuring Unrestricted Proxy ARP (CLI Procedure) on page 1153

port-error-disable

Syntax

```
port-error-disable {
  disable-timeout timeout ;
}
```

Hierarchy Level [edit **ethernet-switching-options**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable rather than block an interface when enforcing MAC limiting, MAC move limiting, and rate-limiting configuration options for shutting down the interface; and allow the interface to recovery automatically from the error condition after a specified period of time:

NOTE: The **port-error-disable** configuration does not apply to pre-existing error conditions. It impacts only error conditions that are detected after **port-error-disable** has been enabled and committed. To clear a pre-existing error condition and restore the interface to service, use the clear ethernet-switching port-error command.

- If you have enabled **mac-limit** with the **shutdown** option and enable **port-error-disable**, the switch disables (rather than shuts down) the interface when the MAC address limit is reached.
- If you have enabled **mac-move-limit** with the **shutdown** option and you enable **port-error-disable**, the switch disables (rather than shuts down) the interface when the maximum number of moves to a new interface is reached.
- If you have enabled **storm-control** with the **action-shutdown** option and you enable **port-error-disable**, the switch disables (rather than shuts down) the interface when broadcast traffic and unknown unicast traffic exceeds the specified levels.

Default Not enabled.

Required Privilege Level routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation

- Configuring Autorecovery From the Disabled State on Secure or Storm Control Interfaces (CLI Procedure) on page 2516
- Configuring Port Security (CLI Procedure) on page 2626
- Example: Configuring Storm Control to Prevent Network Outages on J-EX Series Switches on page 2513

prefix

Syntax prefix hostname;

Hierarchy Level [edit **ethernet-switching-options secure-access-port vlan** (all | *vlan-name*) **dhcp-option82 circuit-id**]
[edit forwarding-options helpers bootp **dhcp-option82 circuit-id**]
[edit forwarding-options helpers bootp interface *interface-name* **dhcp-option82 circuit-id**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure an optional prefix for the circuit ID suboption in the DHCP option 82 information that is inserted by the switch into the packet header of a DHCP request before it forwards or relays the request to a DHCP server.

Default If **prefix** is not explicitly specified, no prefix is appended to the circuit ID. When **prefix** is specified, it is specified as **prefix hostname** (and the value is the hostname of the switch).

Options **hostname**—Name of the host system (the switch) that is forwarding or relaying the DHCP request from the DHCP client to the DHCP server.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Setting Up DHCP Option 82 on a J-EX Series Switch with No Relay Agent Between Clients and DHCP Server on page 2617
- Example: Setting Up DHCP Option 82 with a J-EX Series Switch as Relay Agent Between Clients and a DHCP Server on page 2615
- Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 2649
- Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 2646
- [edit forwarding-options] Configuration Statement Hierarchy on page 43
- RFC 3046, *DHCP Relay Agent Information Option*, at http://tools.ietf.org/html/rfc3046.

prefix

Syntax prefix hostname | mac | none;

Hierarchy Level [edit **ethernet-switching-options secure-access-port vlan** (all | *vlan-name*) **dhcp-option82 remote-id**]
[edit **forwarding-options helpers bootp dhcp-option82 remote-id**]
[edit **forwarding-options helpers bootp interface** *interface-name* **dhcp-option82 remote-id**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure an optional prefix for the remote ID suboption in the DHCP option 82 information that is inserted by the switch into the packet header of a DHCP request before it forwards or relays the request to a DHCP server.

Default If **prefix** is not explicitly specified, no prefix is appended to the remote ID.

Options

hostname—Name of the host system (the switch) that is forwarding or relaying the DHCP request from the DHCP client to the DHCP server.

mac—MAC address of the host system (the switch) that is forwarding or relaying the DHCP request from the DHCP client to the DHCP server.

none—No prefix is applied to the remote ID.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Setting Up DHCP Option 82 on a J-EX Series Switch with No Relay Agent Between Clients and DHCP Server on page 2617
- Example: Setting Up DHCP Option 82 with a J-EX Series Switch as Relay Agent Between Clients and a DHCP Server on page 2615
- Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 2649
- Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 2646
- [edit forwarding-options] Configuration Statement Hierarchy on page 43
- RFC 3046, *DHCP Relay Agent Information Option*, at http://tools.ietf.org/html/rfc3046.

proxy-arp

Syntax proxy-arp <restricted | unrestricted>;

Hierarchy Level [edit **interfaces** *interface-name* **unit** *logical-unit-number*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the switch to respond to an ARP request if the switch has an active route to the ARP request's target (destination).

Default Proxy ARP is not enabled. The switch responds to an ARP request only if the destination IP address is its own.

Options none—The switch responds to any ARP request for a local or remote address if the switch has a route to the target IP address.

restricted—(Optional) The switch responds to ARP requests in which the physical networks of the source and target are different, and does not respond if the source and target IP addresses are in the same subnet. The switch must also have a route to the target IP address.

unrestricted—(Optional) The switch responds to any ARP request for a local or remote address if the switch has a route to the target IP address.

Default: unrestricted

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Proxy ARP on a J-EX Series Switch on page 2621
- Configuring Proxy ARP (CLI Procedure) on page 1153

remote-id

Syntax

```
remote-id {
  prefix hostname | mac | none;
  use-interface-description;
  use-string string;
}
```

Hierarchy Level

[edit **ethernet-switching-options secure-access-port vlan** (all | *vlan-name*) **dhcp-option82**]
[edit **forwarding-options helpers bootp dhcp-option82**]
[edit **forwarding-options helpers bootp interface** *interface-name* **dhcp-option82**]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Insert the **remote-id** suboption of DHCP option 82 (also known as the DHCP relay agent information option) in DHCP request packet headers before forwarding or relaying requests to a DHCP server. This suboption provides a trusted identifier for the host system that has forwarded or relayed requests to the server.

The remaining statements are explained separately.

Default

If **remote-id** is not explicitly set, no remote ID value is inserted in the DHCP request packet header. If the **remote-id** option is specified but is not qualified by a keyword, the MAC address of the host device (the switch) is used as the remote ID.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Setting Up DHCP Option 82 on a J-EX Series Switch with No Relay Agent Between Clients and DHCP Server on page 2617
- Example: Setting Up DHCP Option 82 with a J-EX Series Switch as Relay Agent Between Clients and a DHCP Server on page 2615
- Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 2649
- Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 2646
- [edit forwarding-options] Configuration Statement Hierarchy on page 43
- RFC 3046, *DHCP Relay Agent Information Option*, at http://tools.ietf.org/html/rfc3046.

secure-access-port

Syntax

```
secure-access-port {
  dhcp-snooping-file {
    location local_pathname | remote_URL;
    timeout seconds;
    write-interval seconds;
  }
  interface (all | interface-name) {
    allowed-mac {
      mac-address-list;
    }
    (dhcp-trusted | no-dhcp-trusted);
    mac-limit limit action action;
    no-allowed-mac-log;
    static-ip ip-address {
      vlan  vlan-name;
      mac mac-address;
    }
  }
  vlan (all | vlan-name) {
    (arp-inspection | no-arp-inspection);
    dhcp-option82 {
      circuit-id {
        prefix hostname;
        use-interface-description;
        use-vlan-id;
      }
      remote-id {
        prefix hostname | mac | none;
        use-interface-description;
        use-string string;
      }
      vendor-id <string>;
    }
    (examine-dhcp | no-examine-dhcp);
    (ip-source-guard | no-ip-source-guard);
    mac-move-limit limit action action;
  }
}
```

Hierarchy Level [edit ethernet-switching-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure port security features, including MAC limiting and whether interfaces can receive DHCP responses, and apply dynamic ARP inspection, DHCP snooping, IP source guard, DHCP option 82, and MAC move limiting to no VLANs, specific VLANs, or all VLANs.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Example: Configuring DHCP Snooping, DAI , and MAC Limiting on a J-EX Series Switch with Access to a DHCP Server Through a Second Switch on page 2593
- Example: Configuring IP Source Guard on a Data VLAN That Shares an Interface with a Voice VLAN on page 2608
- Example: Setting Up DHCP Option 82 on a J-EX Series Switch with No Relay Agent Between Clients and DHCP Server on page 2617
- Understanding How to Protect Access Ports on J-EX Series Switches from Common Attacks on page 2546
- Understanding DHCP Snooping for Port Security on J-EX Series Switches on page 2549

static-ip

Syntax

```
static-ip ip-address {
  vlan vlan-name;
  mac mac-address;
}
```

Hierarchy Level [edit **ethernet-switching-options secure-access-port interface** (all|*interface-name*)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Static (fixed) IP address and static MAC address, with an associated VLAN, added to the DHCP snooping database.

Options ***ip-address***—IP address assigned to a device connected on the specified interface.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Static IP Addresses for DHCP Bindings on Access Ports (CLI Procedure) on page 2645

timeout

Syntax timeout *seconds*;

Hierarchy Level [edit **ethernet-switching-options secure-access-port dhcp-snooping-file**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify a timeout value for remote read and write operations. This value determines the amount of time that the switch waits for a remote system to respond when the DHCP snooping database is stored on a remote FTP site.

Default None

Options ***seconds*** —Value in seconds.

Range: 10 through 3600

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Understanding DHCP Snooping for Port Security on J-EX Series Switches on page 2549

traceoptions

Syntax

```
traceoptions {
  file filename <files number> <no-stamp> <replace> <size size> <world-readable |
    no-world-readable>;
  flag flag <disable>;
}
```

Hierarchy Level [edit **ethernet-switching-options**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define global tracing operations for access security features on Ethernet switches.

Default The **traceoptions** feature is disabled by default.

Options

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file *filename* —Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory **/var/log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached (***x*k to specify KB, ***x*m to specify MB, or ***x*g to specify gigabytes), at which point the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000

Default: 3 files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements. You can include the following flags:

- **access-security**—Trace access security events.
- **all**—All tracing operations.
- **config-internals**—Trace internal configuration operations.
- **forwarding-database**—Trace forwarding database and next-hop events.
- **general**—Trace general events.
- **interface**—Trace interface events.
- **ip-source-guard**—Trace IP source guard events.
- **krt**—Trace communications over routing sockets.
- **lib**—Trace library calls.
- **normal**—Trace normal events.

- **parse**—Trace reading of the configuration.
- **regex-parse**—Trace regular-expression parsing operations.
- **rtg**—Trace redundant trunk group events.
- **state**—Trace state transitions.
- **stp**—Trace spanning-tree events.
- **task**—Trace Ethernet-switching task processing.
- **timer**—Trace Ethernet-switching timer processing.
- **vlan**—Trace VLAN events.

no-stamp—(Optional) Do not timestamp the trace file.

Default: If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Restrict file access to the user who created the file.

replace—(Optional) Replace an existing trace file if there is one rather than appending to it.

Default: If you do not include this option, tracing output is appended to an existing trace file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes. When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten. If you specify a maximum number of files, you also must specify a maximum file size with the **files** option.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify gigabytes

Range: 10 KB through 1 gigabyte

Default: 128 KB

world-readable—(Optional) Enable unrestricted file access.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Port Security for J-EX Series Switches Overview on page 2545
- J-EX Series Switches Interfaces Overview on page 863
- Understanding IP Source Guard for Port Security on J-EX Series Switches on page 2563
- Understanding Redundant Trunk Links on J-EX Series Switches on page 1049
- Understanding STP for J-EX Series Switches on page 1275
- Understanding Bridging and VLANs on J-EX Series Switches on page 1041

use-interface-description

Syntax use-interface-description;

Hierarchy Level
[edit **ethernet-switching-options secure-access-port vlan** (all | *vlan-name*) **dhcp-option82 circuit-id**]
[edit forwarding-options helpers bootp **dhcp-option82 circuit-id**]
[edit forwarding-options helpers bootp interface *interface-name* **dhcp-option82 circuit-id**]

[edit **ethernet-switching-options secure-access-port vlan** (all | *vlan-name*) **dhcp-option82 remote-id**]
[edit forwarding-options helpers bootp **dhcp-option82 remote-id**]
[edit forwarding-options helpers bootp interface *interface-name* **dhcp-option82 remote-id**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Use the interface description rather than the interface name (the default) in the circuit ID or remote ID value in the DHCP option 82 information.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Setting Up DHCP Option 82 on a J-EX Series Switch with No Relay Agent Between Clients and DHCP Server on page 2617
- Example: Setting Up DHCP Option 82 with a J-EX Series Switch as Relay Agent Between Clients and a DHCP Server on page 2615
- Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 2649
- Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 2646
- [edit forwarding-options] Configuration Statement Hierarchy on page 43
- RFC 3046, *DHCP Relay Agent Information Option*, at http://tools.ietf.org/html/rfc3046.

use-string

Syntax use-string *string*;

Hierarchy Level [edit **ethernet-switching-options secure-access-port vlan** (all | *vlan-name*) **dhcp-option82 remote-id**]
[edit forwarding-options helpers bootp **dhcp-option82 remote-id**]
[edit forwarding-options helpers bootp interface *interface-name* **dhcp-option82 remote-id**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Use a string rather than the MAC address of the host system (the default) in the remote ID value in the DHCP option 82 information.

Options ***string***—Character string used as the remote ID value.

Range:1–255 characters

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Setting Up DHCP Option 82 on a J-EX Series Switch with No Relay Agent Between Clients and DHCP Server on page 2617
- Example: Setting Up DHCP Option 82 with a J-EX Series Switch as Relay Agent Between Clients and a DHCP Server on page 2615
- Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 2649
- Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 2646
- [edit forwarding-options] Configuration Statement Hierarchy on page 43
- RFC 3046, *DHCP Relay Agent Information Option*, at http://tools.ietf.org/html/rfc3046.

use-vlan-id

Syntax use-vlan-id;

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Use the VLAN ID rather than the VLAN name (the default) in the circuit ID value in the DHCP option 82 information.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Setting Up DHCP Option 82 on a J-EX Series Switch with No Relay Agent Between Clients and DHCP Server on page 2617
- Example: Setting Up DHCP Option 82 with a J-EX Series Switch as Relay Agent Between Clients and a DHCP Server on page 2615
- Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 2649
- Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 2646
- [edit forwarding-options] Configuration Statement Hierarchy on page 43
- RFC 3046, *DHCP Relay Agent Information Option*, at http://tools.ietf.org/html/rfc3046.

vendor-id

Syntax vendor-id <*string*>;

Hierarchy Level [edit **ethernet-switching-options secure-access-port vlan** (all | *vlan-name*) **dhcp-option82**]
[edit forwarding-options helpers bootp **dhcp-option82**]
[edit forwarding-options helpers bootp interface *interface-name* **dhcp-option82**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Insert a vendor ID in the DHCP option 82 information in a DHCP request packet header before forwarding or relaying the request to a DHCP server.

Default If **vendor-id** is not explicitly configured for DHCP option 82, no vendor ID is set.

Options ***string***—(Optional) A single string that designates the vendor ID.

Range: 1–255 characters

Default: If you specify **vendor-id** with no ***string*** value, the default vendor ID **Juniper** is configured.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Setting Up DHCP Option 82 on a J-EX Series Switch with No Relay Agent Between Clients and DHCP Server on page 2617
- Example: Setting Up DHCP Option 82 with a J-EX Series Switch as Relay Agent Between Clients and a DHCP Server on page 2615
- Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 2649
- Setting Up DHCP Option 82 with the Switch as a Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 2646
- [edit forwarding-options] Configuration Statement Hierarchy on page 43

vlan

Syntax

```
vlan (all | vlan-name) {
  (arp-inspection | no-arp-inspection);
  dhcp-option82 {
    circuit-id {
      prefix hostname;
      use-interface-description;
      use-vlan-id;
    }
    remote-id {
      prefix hostname | mac | none;
      use-interface-description;
      use-string string;
    }
    vendor-id <string>;
  }
  (examine-dhcp | no-examine-dhcp);
  (ip-source-guard | no-ip-source-guard);
  mac-move-limit limit action action;
}
```

Hierarchy Level [edit ethernet-switching-options secure-access-port]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply DHCP snooping, dynamic ARP inspection (DAI), IP source guard, DHCP option 82, and MAC move limiting.

The remaining statements are explained separately.

TIP: To display a list of all configured VLANs on the system, including VLANs that are configured but not committed, type ? after vlan or vlans in your configuration mode command line. Note that only one VLAN is displayed for a VLAN range.

Options

all—Apply DHCP snooping, DAI, IP source guard, DHCP option 82, and MAC move limiting to all VLANs.

vlan-name—Apply DHCP snooping, DAI, IP source guard, DHCP option 82, and MAC move limiting to the specified VLAN.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Example: Configuring IP Source Guard with Other J-EX Series Switch Features to Mitigate Address-Spoofing Attacks on Untrusted Access Interfaces on page 2600

- Example: Setting Up DHCP Option 82 on a J-EX Series Switch with No Relay Agent Between Clients and DHCP Server on page 2617
- Enabling Dynamic ARP Inspection (CLI Procedure) on page 2633
- Enabling DHCP Snooping (CLI Procedure) on page 2630
- Configuring IP Source Guard (CLI Procedure) on page 2643
- Configuring MAC Move Limiting (CLI Procedure) on page 2639
- Setting Up DHCP Option 82 on the Switch with No Relay Agent Between Clients and DHCP Server (CLI Procedure) on page 2649

vlan

Syntax vlan *vlan-name*;

Hierarchy Level [edit **ethernet-switching-options secure-access-port interface** (all | *interface-name*) **static-ip** *ip-address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Associate the static IP address with the specified VLAN associated with the specified interface.

Options ***vlan-name*** —Name of a specific VLAN associated with the specified interface.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Static IP Addresses for DHCP Bindings on Access Ports (CLI Procedure) on page 2645

write-interval

Syntax write-interval *seconds*;

Hierarchy Level [edit **ethernet-switching-options secure-access-port dhcp-snooping-file**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify how frequently the switch writes the database entries from memory into the specified DHCP snooping database file.

Default None

Options ***seconds***—Value in seconds.
Range: 60 through 86400

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Understanding DHCP Snooping for Port Security on J-EX Series Switches on page 2549

CHAPTER 99

Operational Mode Commands for Port Security

clear arp inspection statistics

Syntax clear arp inspection statistics
<interface *interface*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Clear ARP inspection statistics.

Options none—Clears ARP statistics on all interfaces.

interface *interface-names*—(Optional) Clear ARP statistics on one or more interfaces.

Required Privilege Level clear

Related Documentation

- **show arp inspection statistics on page 2709**
- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Verifying That DAI Is Working Correctly on page 2656

List of Sample Output **clear arp inspection statistics on page 2706**

Output Fields This command produces no output.

clear arp inspection statistics

```
user@switch> clear arp inspection statistics
```

clear dhcp snooping binding

Syntax

clear dhcp snooping binding
<mac (all | *mac-address*)>
<vlan (all | *vlan-name*)>
<vlan (all | *vlan-name*) mac (all | *mac-address*)>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Clear the DHCP snooping database information.

Options

mac (all | ***mac-address***)—(Optional) Clear DHCP snooping information for the specified MAC address or all MAC addresses.

vlan (all | ***vlan-name***)—(Optional) Clear DHCP snooping information for the specified VLAN or all VLANs.

Required Privilege Level

clear

Related Documentation

- **show dhcp snooping binding on page 2710**
- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Verifying That DHCP Snooping Is Working Correctly on page 2654

List of Sample Output

clear dhcp snooping binding on page 2707

Output Fields

This command produces no output.

clear dhcp snooping binding

```
user@switch> clear dhcp snooping binding
```

clear dhcp snooping statistics

Syntax **clear dhcp snooping statistics**

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Clear all Dynamic Host Configuration Protocol (DHCP) snooping statistics.

Required Privilege Level view

Related Documentation

- **show dhcp snooping statistics on page 2711**
- Understanding DHCP Snooping for Port Security on J-EX Series Switches on page 2549

List of Sample Output **clear dhcp snooping statistics on page 2708**

Output Fields See **show dhcp snooping statistics** for an explanation of the output fields.

clear dhcp snooping statistics The following sample output displays the DHCP snooping statistics before and after the **clear dhcp snooping statistics** command is issued.

```
user@switch> show dhcp snooping statistics
Successful Transfers :        0   Failed Transfers :       21
Successful Reads     :        0   Failed Reads     :        0
Successful Writes    :        0   Failed Writes    :       21

user@switch> clear dhcp snooping statistics

user@switch> show dhcp snooping statistics
Successful Transfers :        0   Failed Transfers :        0
Successful Reads     :        0   Failed Reads     :        0
Successful Writes    :        0   Failed Writes    :        0
```

show arp inspection statistics

Syntax show arp inspection statistics

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display ARP inspection statistics.

Required Privilege Level view

Related Documentation

- **clear arp inspection statistics on page 2706**
- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Verifying That DAI Is Working Correctly on page 2656

List of Sample Output **show arp inspection statistics on page 2709**

Output Fields Table 330 on page 2709 lists the output fields for the **show arp inspection statistics** command. Output fields are listed in the approximate order in which they appear.

Table 330: show arp inspection statistics Output Fields

Field Name	Field Description	Level of Output
Interface	Interface on which ARP inspection has been applied.	All levels
Packets received	Total number of packets total that underwent ARP inspection.	All levels
ARP inspection pass	Total number of packets that passed ARP inspection.	All levels
ARP inspection failed	Total number of packets that failed ARP inspection.	All levels

show arp inspection statistics

```
user@switch> show arp inspection statistics

Interface     Packets received     ARP inspection pass  ARP inspection failed
---------     ----------------     -------------------  ---------------------
  ge-0/0/0                   0                       0                     0
  ge-0/0/1                   0                       0                     0
  ge-0/0/2                   0                       0                     0
  ge-0/0/3                   0                       0                     0
  ge-0/0/4                   0                       0                     0
  ge-0/0/5                   0                       0                     0
  ge-0/0/6                   0                       0                     0
  ge-0/0/7                 703                     701                     2
```

show dhcp snooping binding

Syntax show dhcp snooping binding

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display the DHCP snooping database information.

Required Privilege Level view

Related Documentation

- **clear dhcp snooping binding on page 2707**
- Example: Configuring Port Security, with DHCP Snooping, DAI, MAC Limiting, and MAC Move Limiting, on a J-EX Series Switch on page 2569
- Verifying That DHCP Snooping Is Working Correctly on page 2654

List of Sample Output **show dhcp snooping binding on page 2710**

Output Fields Table 331 on page 2710 lists the output fields for the **show dhcp snooping binding** command. Output fields are listed in the approximate order in which they appear.

Table 331: show dhcp snooping binding Output Fields

Field Name	Field Description	Level of Output
MAC Address	MAC address of the network device; bound to the IP address.	All levels
IP Address	IP address of the network device; bound to the MAC address.	All levels
Lease	Lease granted to the IP address.	All levels
Type	How the MAC address was acquired.	All levels
VLAN	VLAN name of the network device whose MAC address is shown.	All levels
Interface	Interface address (port).	All levels

show dhcp snooping binding

```
user@switch> show dhcp snooping binding

DHCP Snooping Information:
MAC Address             IP Address Lease   Type     VLAN    Interface
-----------------       ---------- -----   -------  ----    ---------
00:00:01:00:00:03       192.0.2.0  640     dynamic  guest    ge-0/0/12.0
00:00:01:00:00:04       192.0.2.1  720     dynamic  guest    ge-0/0/12.0
00:00:01:00:00:05       192.0.2.5  800     dynamic  guest    ge-0/0/13.0
```

show dhcp snooping statistics

Syntax **show dhcp snooping statistics**

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display statistics for read and write operations to the DHCP snooping database.

Required Privilege Level view

Related Documentation
- **clear dhcp snooping statistics on page 2708**
- Understanding DHCP Snooping for Port Security on J-EX Series Switches on page 2549

List of Sample Output **show dhcp snooping statistics on page 2711**

Output Fields Table 332 on page 2711 lists the output fields for the **show dhcp snooping statistics** command. Output fields are listed in the approximate order in which they appear.

Table 332: show dhcp snooping statistics Output Fields

Field Name	Field Description
Successful Transfers	Number of entries successfully transferred from memory to the DHCP snooping database.
Successful Reads	Number of entries successfully read from memory to the DHCP snooping database.
Successful Writes	Number of entries successfully written from memory to the DHCP snooping database.
Failed Transfers	Number of entries that failed being transferred from memory to the DHCP snooping database.
Failed Reads	Number of entries that failed being read from memory to the DHCP snooping database.
Failed Writes	Number of entries that failed being written from memory to the DHCP snooping database.

show dhcp snooping statistics

```
user@switch> show dhcp snooping statistics
Successful Transfers :        0   Failed Transfers :       21
Successful Reads     :        0   Failed Reads     :        0
Successful Writes    :        0   Failed Writes    :       21
```

show ethernet-switching table

Syntax

show ethernet-switching table
<brief | detail | extensive | summary>
<interface *interface-name*>
<management-vlan>
<sort-by (*name* | *tag*)
<vlan (*vlan-name*)>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Displays the Ethernet switching table.

Options

none—(Optional) Display brief information about the Ethernet switching table.

brief | detail | extensive | summary—(Optional) Display the specified level of output.

management-vlan—(Optional) Display the Ethernet switching table for a management VLAN.

interface-name—(Optional) Display the Ethernet switching table for a specific interface.

sort-by (*name* | *tag*)—(Optional) Display VLANs in ascending order of VLAN IDs or VLAN names.

vlan *vlan-name*—(Optional) Display the Ethernet switching table for a specific VLAN.

Required Privilege Level

view

Related Documentation

- Example: Setting Up Basic Bridging and a VLAN for a J-EX Series Switch on page 1063
- Example: Setting Up Bridging with Multiple VLANs for J-EX Series Switches on page 1070
- Example: Configure Automatic VLAN Administration Using GVRP on page 1087
- Example: Setting Up Q-in-Q Tunneling on J-EX Series Switches on page 1105

List of Sample Output

show ethernet-switching table on page 2713
show ethernet-switching table brief on page 2714
show ethernet-switching table detail on page 2714
show ethernet-switching table extensive on page 2715
show ethernet-switching table interface ge-0/0/1 on page 2715

Output Fields

Table 333 on page 2712 lists the output fields for the **show ethernet-switching table** command. Output fields are listed in the approximate order in which they appear.

Table 333: show ethernet-switching table Output Fields

Field Name	Field Description	Level of Output
VLAN	The name of a VLAN.	All levels

Table 333: show ethernet-switching table Output Fields *(continued)*

Field Name	Field Description	Level of Output
Tag	The VLAN ID tag name or number.	**extensive**
MAC or **MAC address**	The MAC address associated with the VLAN.	All levels
Type	The type of MAC address. Values are: • **static**—The MAC address is manually created. • **learn**—The MAC address is learned dynamically from a packet's source MAC address. • **flood**—The MAC address is unknown and flooded to all members.	All levels
Age	The time remaining before the entry ages out and is removed from the Ethernet switching table.	All levels
Interfaces	Interface associated with learned MAC addresses or **All-members** (flood entry).	All levels
Learned	For learned entries, the time which the entry was added to the Ethernet-switching table.	**detail**, **extensive**
Nexthop index	The nexthop index number.	**detail**, **extensive**

show ethernet-switching table

```
user@switch> show ethernet-switching table
Ethernet-switching table: 57 entries, 17 learned
  VLAN              MAC address       Type         Age Interfaces
  F2                *                 Flood          - All-members
  F2                00:00:05:00:00:03 Learn          0 ge-0/0/44.0
  F2                00:19:e2:50:7d:e0 Static         - Router
  Linux             *                 Flood          - All-members
  Linux             00:19:e2:50:7d:e0 Static         - Router
  Linux             00:30:48:90:54:89 Learn          0 ge-0/0/47.0
  T1                *                 Flood          - All-members
  T1                00:00:05:00:00:01 Learn          0 ge-0/0/46.0
  T1                00:00:5e:00:01:00 Static         - Router
  T1                00:19:e2:50:63:e0 Learn          0 ge-0/0/46.0
  T1                00:19:e2:50:7d:e0 Static         - Router
  T10               *                 Flood          - All-members
  T10               00:00:5e:00:01:09 Static         - Router
  T10               00:19:e2:50:63:e0 Learn          0 ge-0/0/46.0
  T10               00:19:e2:50:7d:e0 Static         - Router
  T111              *                 Flood          - All-members
  T111              00:19:e2:50:63:e0 Learn          0 ge-0/0/15.0
  T111              00:19:e2:50:7d:e0 Static         - Router
  T111              00:19:e2:50:ac:00 Learn          0 ge-0/0/15.0
  T2                *                 Flood          - All-members
  T2                00:00:5e:00:01:01 Static         - Router
  T2                00:19:e2:50:63:e0 Learn          0 ge-0/0/46.0
  T2                00:19:e2:50:7d:e0 Static         - Router
  T3                *                 Flood          - All-members
  T3                00:00:5e:00:01:02 Static         - Router
  T3                00:19:e2:50:63:e0 Learn          0 ge-0/0/46.0
  T3                00:19:e2:50:7d:e0 Static         - Router
  T4                *                 Flood          - All-members
```

```
  T4                00:00:5e:00:01:03 Static         - Router
  T4                00:19:e2:50:63:e0 Learn          0 ge-0/0/46.0
[output truncated]
```

show ethernet-switching table brief

```
user@switch> show ethernet-switching table brief
Ethernet-switching table: 57 entries, 17 learned
  VLAN              MAC address       Type          Age Interfaces
  F2                *                 Flood           - All-members
  F2                00:00:05:00:00:03 Learn           0 ge-0/0/44.0
  F2                00:19:e2:50:7d:e0 Static          - Router
  Linux             *                 Flood           - All-members
  Linux             00:19:e2:50:7d:e0 Static          - Router
  Linux             00:30:48:90:54:89 Learn           0 ge-0/0/47.0
  T1                *                 Flood           - All-members
  T1                00:00:05:00:00:01 Learn           0 ge-0/0/46.0
  T1                00:00:5e:00:01:00 Static          - Router
  T1                00:19:e2:50:63:e0 Learn           0 ge-0/0/46.0
  T1                00:19:e2:50:7d:e0 Static          - Router
  T10               *                 Flood           - All-members
  T10               00:00:5e:00:01:09 Static          - Router
  T10               00:19:e2:50:63:e0 Learn           0 ge-0/0/46.0
  T10               00:19:e2:50:7d:e0 Static          - Router
  T111              *                 Flood           - All-members
  T111              00:19:e2:50:63:e0 Learn           0 ge-0/0/15.0
  T111              00:19:e2:50:7d:e0 Static          - Router
  T111              00:19:e2:50:ac:00 Learn           0 ge-0/0/15.0
  T2                *                 Flood           - All-members
  T2                00:00:5e:00:01:01 Static          - Router
  T2                00:19:e2:50:63:e0 Learn           0 ge-0/0/46.0
  T2                00:19:e2:50:7d:e0 Static          - Router
  T3                *                 Flood           - All-members
  T3                00:00:5e:00:01:02 Static          - Router
  T3                00:19:e2:50:63:e0 Learn           0 ge-0/0/46.0
  T3                00:19:e2:50:7d:e0 Static          - Router
  T4                *                 Flood           - All-members
  T4                00:00:5e:00:01:03 Static          - Router
  T4                00:19:e2:50:63:e0 Learn           0 ge-0/0/46.0
[output truncated]
```

show ethernet-switching table detail

```
user@switch> show ethernet-switching table detail
Ethernet-switching table: 5 entries, 2 learned
 VLAN: default, Tag: 0, MAC: *, Interface: All-members
  Interfaces:
   ge-0/0/11.0, ge-0/0/20.0, ge-0/0/30.0, ge-0/0/36.0, ge-0/0/3.0
  Type: Flood
  Nexthop index: 1307

 VLAN: default, Tag: 0, MAC: 00:1f:12:30:b8:83, Interface: ge-0/0/3.0
   Type: Learn, Age: 0, Learned: 20:09:26
   Nexthop index: 1315

 VLAN: v1, Tag: 101, MAC: *, Interface: All-members
  Interfaces:
   ge-0/0/31.0
  Type: Flood
  Nexthop index: 1313

 VLAN: v1, Tag: 101, MAC: 00:1f:12:30:b8:89, Interface: ge-0/0/31.0
  Type: Learn, Age: 0, Learned: 20:09:25
  Nexthop index: 1312
```

```
 VLAN: v2, Tag: 102, MAC: *, Interface: All-members
  Interfaces:
   ae0.0
  Type: Flood
  Nexthop index: 1317
```

show ethernet-switching table extensive

```
user@switch> show ethernet-switching table extensive
Ethernet-switching table: 3 entries, 1 learned

    VLAN: v1, Tag: 10, MAC: *, Interface: All-members
    Interfaces:
                ge-0/0/14.0, ge-0/0/1.0, ge-0/0/2.0, ge-0/0/3.0, ge-0/0/4.0,
                ge-0/0/5.0, ge-0/0/6.0, ge-0/0/7.0, ge-0/0/8.0, ge-0/0/10.0,
                ge-0/0/0.0
    Type: Flood
    Nexthop index: 567

    VLAN: v1, Tag: 10, MAC: 00:21:59:c6:93:22, Interface: Router
    Type: Static
    Nexthop index: 0

    VLAN: v1, Tag: 10, MAC: 00:21:59:c9:9a:4e, Interface: ge-0/0/14.0
    Type: Learn, Age: 0, Learned: 18:40:50
    Nexthop index: 564
```

show ethernet-switching table interface ge-0/0/1

```
user@switch> show ethernet-switching table interface ge-0/0/1
Ethernet-switching table: 1 unicast entries
  VLAN              MAC address       Type         Age Interfaces
  V1                *                 Flood          - All-members
  V1                00:00:05:00:00:05 Learn          0 ge-0/0/1.0
```

show ip-source-guard

Syntax **show ip-source-guard**

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display IP source guard database information.

Required Privilege Level view

Related Documentation

- Example: Configuring IP Source Guard on a Data VLAN That Shares an Interface with a Voice VLAN on page 2608
- Example: Configuring IP Source Guard with Other J-EX Series Switch Features to Mitigate Address-Spoofing Attacks on Untrusted Access Interfaces on page 2600
- Verifying That IP Source Guard Is Working Correctly on page 2662

List of Sample Output **show ip-source-guard on page 2716**

Output Fields Table 334 on page 2716 lists the output fields for the **show ip-source-guard** command. Output fields are listed in the approximate order in which they appear.

Table 334: show ip-source-guard Output Fields

Field Name	Field Description
VLAN	VLAN on which IP source guard is enabled.
Interface	Access interface associated with the VLAN in column 1.
Tag	VLAN ID for the VLAN in column 1. Possible values are: • 0, indicating the VLAN is not tagged. • 1 – 4093
IP Address	Source IP address for a device connected to the interface in column 2. A value of * (star, or asterisk) indicates that IP source guard is not enabled on this VLAN but the interface is shared with a VLAN that is enabled for IP source guard.
MAC Address	Source MAC address for a device connected to the interface in column 2. A value of * (star, or asterisk) indicates that IP source guard is not enabled on this VLAN but the interface is shared with a VLAN that is enabled for IP source guard.

show ip-source-guard

```
user@switch> show ip-source-guard
IP source guard information:
Interface    Tag  IP Address   MAC Address        VLAN

ge-0/0/12.0  0    10.10.10.7   00:30:48:92:A5:9D  vlan100

ge-0/0/13.0  0    10.10.10.9   00:30:48:8D:01:3D  vlan100
```

```
ge-0/0/13.0  100  *              *                   voice
```

show system statistics arp

Syntax show system statistics arp

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display system-wide Address Resolution Protocol (ARP) statistics.

Required Privilege Level view

Related Documentation
- Example: Configuring Unrestricted Proxy ARP on a J-EX Series Switch on page 2621
- Verifying That Unrestricted Proxy ARP Is Working Correctly on page 1164

Sample Output

```
user@switch> show system statistics arp
arp:
        90060 datagrams received
        34 ARP requests received
        610 ARP replies received
        0 resolution request  received
        0 unrestricted proxy requests
        0 restricted proxy requests
        0 received proxy requests
        0 unrestricted proxy requests not proxied
        0 restricted proxy requests not proxied
        0 datagrams with bogus interface
        0 datagrams with incorrect length
        0 datagrams for non-IP protocol
        0 datagrams with unsupported op code
        0 datagrams with bad protocol address length
        0 datagrams with bad hardware address length
        0 datagrams with multicast source address
        0 datagrams with multicast source address
        0 datagrams with my own hardware address
        0 datagrams for an address not on the interface
        0 datagrams with a broadcast source address
        294 datagrams with source address duplicate to mine
        89113 datagrams which were not for me
        0 packets discarded waiting for resolution
        0 packets sent after waiting for resolution
        309 ARP requests sent
        35 ARP replies sent
        0 requests for memory denied
        0 requests dropped on entry
        0 requests dropped during retry
        0 requests dropped due to interface deletion
        0 requests on unnumbered interfaces
        0 new requests on unnumbered interfaces
        0 replies for from unnumbered interfaces
        0 requests on unnumbered interface with non-subnetted donor
        0 replies from unnumbered interface with non-subnetted donor
```

PART 20

Routing Policy and Packet Filtering (Firewall Filters)

CHAPTER 100

Firewall Filters—Overview

Firewall Filters for J-EX Series Switches Overview

Firewall filters provide rules that define whether to permit, deny, or forward packets that are transiting an interface on a J-EX Series Switch from a source address to a destination address. You configure firewall filters to determine whether to permit, deny, or forward traffic before it enters or exits a port, VLAN, or Layer 3 (routed) interface to which the firewall filter is applied. An *ingress* firewall filter is a filter that is applied to packets that are entering a network. An *egress* firewall filter is a filter that is applied to packets that are exiting a network. You can configure firewall filters to subject packets to filtering, class-of-service (CoS) marking (grouping similar types of traffic together, and treating each type of traffic as a class with its own level of service priority), and traffic policing (controlling the maximum rate of traffic sent or received on an interface).

Firewall Filter Types

The following firewall filter types are supported for J-EX Series switches:

- Port (Layer 2) firewall filter—Port firewall filters apply to Layer 2 switch ports. You can apply port firewall filters in both ingress and egress directions on a physical port.

- VLAN firewall filter—VLAN firewall filters provide access control for packets that enter a VLAN, are bridged within a VLAN, and leave a VLAN. You can apply VLAN firewall filters in both ingress and egress directions on a VLAN. VLAN firewall filters are applied to all packets that are forwarded to or forwarded from the VLAN.
- Router (Layer 3) firewall filter—You can apply a router firewall filter in both ingress and egress directions on Layer 3 (routed) interfaces and routed VLAN interfaces (RVIs). You can apply a router firewall filter in the ingress direction on the loopback interface (**lo0**) also.

NOTE: Firewall filters configured on loopback interfaces are applied only to packets that are sent to the routing engine CPU for further processing. Firewall filters are not applied to packets transiting the management interface (me0).

On J-EX4200 and J-EX8200 Ethernet switches, you can apply a router firewall filter to both IPv4 and IPv6 traffic. You can apply firewall filter match conditions to IPv6 traffic on Layer 3 interfaces, aggregated Ethernet interfaces, and loopback interfaces. To configure port firewall filters and VLAN firewall filters for IPv6 traffic, you must include the match condition **ether-type ipv6** and apply the filter on Layer 2 interfaces or VLANs. When you include the match condition **ether-type ipv6** in a term, you must ensure that other match conditions specified in the term are valid for IPv6 traffic. If the port firewall filter or VLAN firewall filter term contains the match condition **ether-type ipv6**, with no other IPv6 match condition specified, all IPv6 traffic is matched.

NOTE: A term without the match condition ether-type ipv6 applies only to IPv4 traffic, and a term with that match condition applies only to IPv6 traffic. Hence, to configure port and VLAN firewall filters for both IPv4 and IPv6 traffic, you should configure two different terms, once each for IPv4 and IPv6 traffic.

To apply a firewall filter, you must:

1. Configure the firewall filter.
2. Apply the firewall filter to a port, VLAN, or Layer 3 interface. You can apply a firewall filter to aggregated Ethernet interfaces and loopback interfaces also.

Firewall Filter Components

In a firewall filter, you first define the family address type, (**ethernet-switching**, **inet**, or **inet6**), and then you define one or more terms that specify the filtering criteria and the action to take if a match occurs.

The maximum number of terms allowed per firewall filter depends on the switch platform:

- 2048 for J-EX4200 switches—as allocated by the dynamic allocation of Ternary Content Addressable Memory (TCAM) for port, VLAN, and router firewall filters

- Determined by the dynamic allocation of TCAM for port, VLAN, and router firewall filters on J-EX8200 switches

NOTE: The on-demand dynamic allocation of the shared space TCAM in J-EX8200 switches is achieved by assigning free space blocks to firewall filters. Firewall filters are categorized into two different pools. Port and VLAN filters are pooled together (the memory threshold for this pool is 22K) while router firewall filters are pooled separately (the threshold for this pool is 32K). The assignment happens based on the filter pool type. Free space blocks can be shared only among the firewall filters belonging to the same filter pool type. An error message is generated when you try to configure a firewall filter beyond the TCAM threshold.

Each term consists of the following components:

- Match conditions—Specifies the values or fields that the packet must contain. You can define various match conditions, including the IP source address field, IP destination address field, Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) source port field, IP protocol field, Internet Control Message Protocol (ICMP) packet type, TCP flags, and interfaces.
- Action—Specifies what to do if a packet matches the match conditions. Possible actions are to accept or discard the packet or to send the packet to a specific virtual routing interface. In addition, packets can be counted to collect statistical information. If no action is specified for a term, the default action is to accept the packet.

Firewall Filter Processing

The order of the terms within a firewall filter configuration is important. Packets are tested against each term in the order in which the terms are listed in the firewall filter configuration. When a firewall filter contains multiple terms, the switch takes a top-down approach and compares a packet against the first term in the firewall filter. If the packet matches the first term, the switch executes the action defined by that term to either permit or deny the packet, and no other terms are evaluated. If the switch does not find a match between the packet and first term, it compares the packet to the next term in the firewall filter by using the same match process. If no match occurs between the packet and the second term, the switch continues to compare the packet to each successive term defined in the firewall filter until a match is found. If a packet does not match any terms in a firewall filter, the default action is to discard the packet.

Related Documentation

- Understanding Planning of Firewall Filters on page 2724
- Understanding Firewall Filter Processing Points for Bridged and Routed Packets on J-EX Series Switches on page 2726
- Understanding How Firewall Filters Are Evaluated on page 2746
- Understanding Firewall Filter Match Conditions on page 2748
- Understanding the Use of Policers in Firewall Filters on page 2752

- Understanding Filter-Based Forwarding for J-EX Series Switches on page 2753
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on J-EX Series Switches on page 2773

Understanding Planning of Firewall Filters

Before you create a firewall filter and apply it to an interface, determine what you want the firewall filter to accomplish and how to use its match conditions and actions to achieve your goals. You must understand how packets are matched to match conditions, the default and configured actions of the firewall filter, and proper placement of the firewall filter.

You can configure and apply no more than one firewall filter per port, VLAN, or router interface, per direction. The following limits apply for the number of firewall filter terms allowed per filter on various switch models:

- On J-EX4200 switches, the number of terms allowed per filter cannot exceed 2048.
- On J-EX8200 switches, the number of terms allowed per filter cannot exceed 32768.

In addition, you should try to be conservative in the number of terms (rules) that you include in each firewall filter because a large number of terms requires longer processing time during a commit and also can make firewall filter testing and troubleshooting more difficult. Similarly, applying firewall filters across many switch and router interfaces can make testing and troubleshooting the rules of those filters difficult.

Before you configure and apply firewall filters, answer the following questions for each of those firewall filters:

1. What is the purpose of the firewall filter?

   For example, you can use a firewall filter to limit traffic to source and destination MAC addresses, specific protocols, or certain data rates or to prevent denial of service (DoS) attacks.

2. What are the appropriate match conditions?

   a. Determine the packet header fields that the packet must contain for a match. Possible fields include:

      - Layer 2 header fields—Source and destination MAC addresses, dot1q tag, Ethernet type, and VLAN
      - Layer 3 header fields—Source and destination IP addresses, protocols, and IP options (IP precedence, IP fragmentation flags, TTL type)
      - TCP header fields—Source and destination ports and flags
      - ICMP header fields—Packet type and code

b. Determine the port, VLAN, or router interface on which the packet was received.

3. What are the appropriate actions to take if a match occurs?

   Possible actions to take if a match occurs are accept, discard, and forward to a routing instance.

4. What additional action modifiers might be required?

   Determine whether additional actions are required if a packet matches a match condition; for example, you can specify an action modifier to count, analyze, or police packets.

5. On what interface should the firewall filter be applied?

   Start with the following basic guidelines:

   - If all the packets entering a port need to be exposed to filtering, then use port firewall filters.
   - If all the packets that are bridged need filtering, then use VLAN firewall filters.
   - If all the packets that are routed need filtering, then use router firewall filters.

   Before you choose the interface on which to apply a firewall filter, understand how that placement can impact traffic flow to other interfaces. In general, apply a firewall filter that filters on source and destination IP addresses, IP protocols, or protocol information—such as ICMP message types, and TCP and UDP port numbers—nearest to the source devices. However, typically apply a firewall filter that filters only on a source IP address nearest to the destination devices. When applied too close to the source device, a firewall filter that filters only on a source IP address could potentially prevent that source device from accessing other services that are available on the network.

   NOTE: Egress firewall filters do not affect the flow of locally generated control packets from the Routing Engine.

6. In which direction should the firewall filter be applied?

   You can apply firewall filters to ports on the switch to filter packets that are entering a port. You can apply firewall filters to VLANs, and Layer 3 (routed) interfaces to filter packets that are entering or exiting a VLAN or routed interface. Typically, you configure different sets of actions for traffic entering an interface than you configure for traffic exiting an interface.

Related Documentation

- Firewall Filters for J-EX Series Switches Overview on page 2721
- Understanding the Use of Policers in Firewall Filters on page 2752
- Understanding How Firewall Filters Are Evaluated on page 2746
- Understanding Filter-Based Forwarding for J-EX Series Switches on page 2753

- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on J-EX Series Switches on page 2773

Understanding Firewall Filter Processing Points for Bridged and Routed Packets on J-EX Series Switches

J-EX Series Switches are multilayered switches that provide Layer 2 switching and Layer 3 routing. You apply firewall filters at multiple processing points in the packet forwarding path on J-EX Series switches. At each processing point, the action to be taken on a packet is determined based on the results of the lookup in the switch's forwarding table. A table lookup determines which exit port on the switch to use to forward the packet.

For both bridged unicast packets and routed unicast packets, firewall filters are evaluated and applied hierarchically. First, a packet is checked against the port firewall filter, if present. If the packet is permitted, it is then checked against the VLAN firewall filter, if present. If the packet is permitted, it is then checked against the router firewall filter, if present. The packet must be permitted by the router firewall filter before it is processed.

Figure 75 on page 2726 shows the various firewall filter processing points in the packet forwarding path in a multilayered switching platform.

Figure 75: Firewall Filter Processing Points in the Packet Forwarding Path

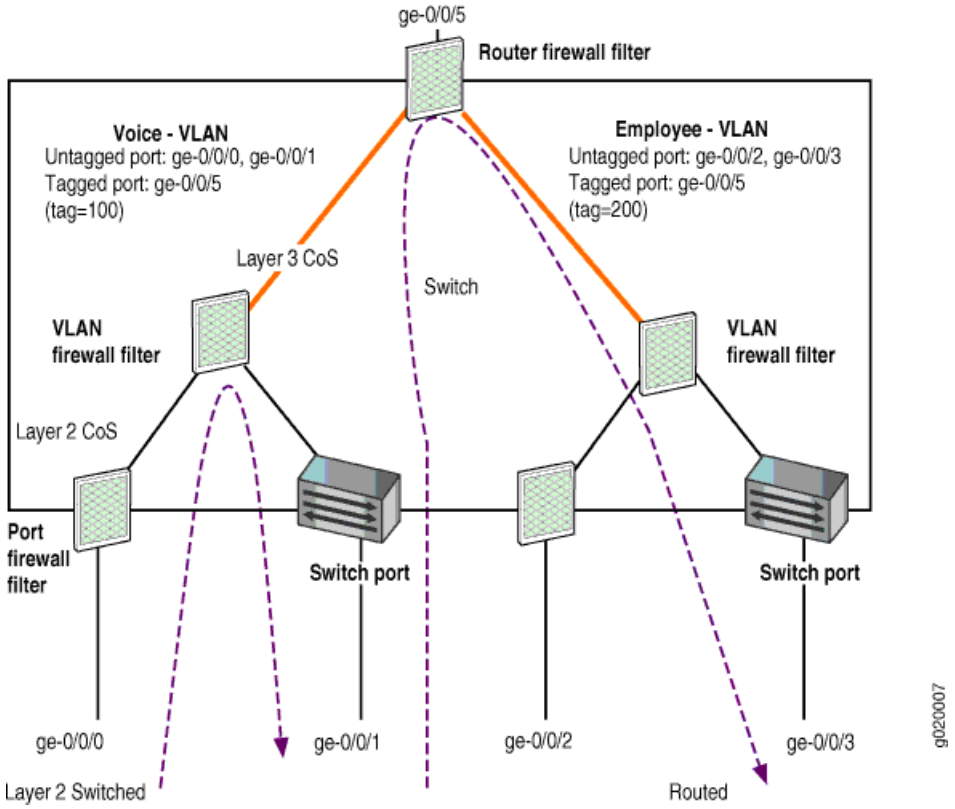

For a multicast packet that results in replications, an egress firewall filter is applied to each copy of the packet based on its corresponding egress VLAN.

For Layer 2 (bridged) unicast packets, the following firewall filter processing points apply:

- Ingress port firewall filter
- Ingress VLAN firewall filter
- Egress port firewall filter
- Egress VLAN firewall filter

For Layer 3 (routed and multilayer-switched) unicast packets, the following firewall filter processing points apply:

- Ingress port firewall filter
- Ingress VLAN firewall filter (Layer 2 CoS)
- Ingress router firewall filter (Layer 3 CoS)
- Egress router firewall filter
- Egress VLAN firewall filter

Related Documentation

- Firewall Filters for J-EX Series Switches Overview on page 2721
- Understanding How Firewall Filters Control Packet Flows on page 2727
- Understanding Bridging and VLANs on J-EX Series Switches on page 1041
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755

Understanding How Firewall Filters Control Packet Flows

J-EX Series Switches support firewall filters that allow you to control flows of data packets and local packets. *Data packets* are chunks of data that transit the switch as they are forwarded from a source to a destination. *Local packets* are chunks of data that are destined for or sent by the switch. Local packets usually contain routing protocol data, data for IP services such as Telnet or SSH, and data for administrative protocols such as the Internet Control Message Protocol (ICMP).

You create firewall filters to protect your switch from excessive traffic transiting the switch to a network destination or destined for the Routing Engine on the switch. Firewall filters that control local packets can also protect your switch from external incidents such as denial-of-service (DoS) attacks.

Firewall filters affect packet flows entering in to or exiting from the switch's interfaces:

- Ingress firewall filters affect the flow of data packets that are received by the switch's interfaces. The Packet Forwarding Engine (PFE) handles this flow. When a switch receives a data packet on an interface, the switch determines where to forward the packet by looking in the forwarding table for the best route (Layer 2 switching, Layer 3

routing) to a destination. Data packets are forwarded to their destination through an outgoing interface. Locally destined packets are forwarded to the Routing Engine.

- Egress firewall filters affect the flow of data packets that are transmitted from the switch's interfaces but do not affect the flow of locally generated control packets from the Routing Engine. The Packet Forwarding Engine handles the flow of data packets that are transmitted from the switch, and egress firewall filters are applied here. The Packet Forwarding Engine also handles the flow of control packets from the Routing Engine.

Figure 76 on page 2728 illustrates the application of ingress and egress firewall filters to control the flow of packets through the switch.

Figure 76: Application of Firewall Filters to Control Packet Flow

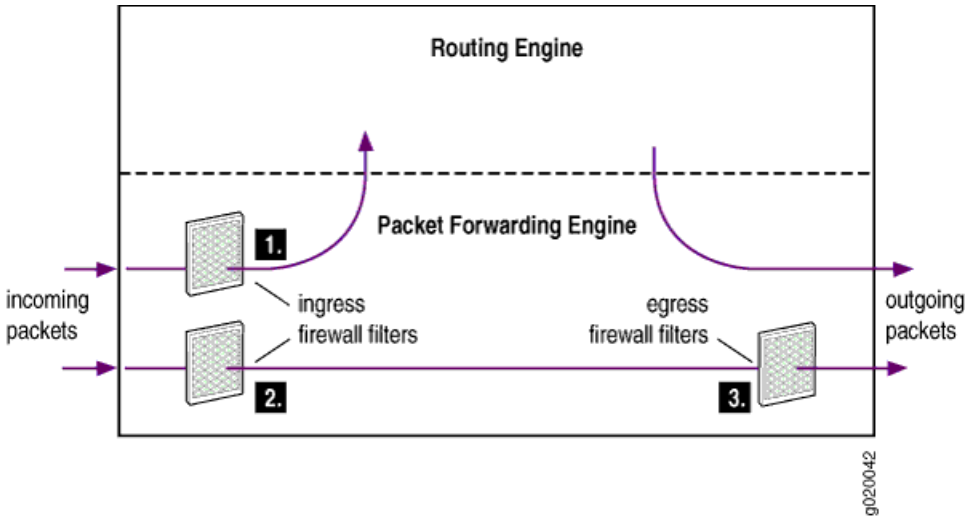

1. Ingress firewall filter applied to control locally destined packets that are received on the switch's interfaces and are destined for the Routing Engine.
2. Ingress firewall filter applied to control incoming packets on the switch's interfaces.
3. Egress firewall filter applied to control packets that are transiting the switch's interfaces.

Related Documentation

- Understanding Firewall Filter Processing Points for Bridged and Routed Packets on J-EX Series Switches on page 2726
- Understanding How Firewall Filters Are Evaluated on page 2746

Firewall Filter Match Conditions and Actions for J-EX Series Switches

Each term in a firewall filter consists of *match conditions* and an *action*. Match conditions are the values or fields that a packet must contain. You can define multiple, single, or no match conditions. If no match conditions are specified for the term, all packets are matched by default. The action is the action that the switch takes if a packet matches the match conditions for the specific term. Action modifiers are optional and specify one

or more actions that the switch takes if a packet matches the match conditions for the specific term. Allowed actions are accept a packet or discard a packet. In addition, you can specify action modifiers to count, mirror, rate limit, and classify packets.

For each firewall filter, you define the terms that specify the filtering criteria (match conditions) to apply to packets and the action for the switch to take if a match occurs.

The string that defines a match condition is called a *match statement*. The following tables list various match conditions and their support platforms, binding points, and actions.

- Table 335 on page 2729 describes the match conditions you can specify when configuring a firewall filter for IPv4 traffic.
- Table 336 on page 2738 describes the match conditions you can specify when configuring a firewall filter for IPv6 traffic.
- Table 337 on page 2744 shows the actions that you can specify in a term.
- Table 338 on page 2744 shows the action modifiers that you can specify in a term.

Table 335: Supported Match Conditions Applicable to IPv4 Traffic for Firewall Filters on J-EX Series Switches

Match Condition	Description	Supported Platforms and Bind Points	
		Ingress	Egress
destination-address *ip-address*	IP destination address field, which is the address of the final destination node. For IPv6, specifies the 128-bit address that is the final destination node address for the packet. The filter description syntax supports the text representations for IPv6 addresses as described in RFC 2373, *IP Version6 Addressing Architecture*.	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—ports, VLANs, and Layer 3 interfaces	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—ports, VLANs, and Layer 3 interfaces
destination-mac-address *mac-address*	Destination media access control (MAC) address of the packet.	• J-EX4200—ports and VLANs • J-EX8200—ports and VLANs	• J-EX4200—ports and VLANs • J-EX8200—ports and VLANs

Table 335: Supported Match Conditions Applicable to IPv4 Traffic for Firewall Filters on J-EX Series Switches *(continued)*

Match Condition	Description	Supported Platforms and Bind Points	
		Ingress	Egress
destination-port *number*	TCP or User Datagram Protocol (UDP) destination port field. Typically, you specify this match in conjunction with the **protocol** match statement to determine which protocol is used on the port. In place of the numeric value, you can specify one of the following text synonyms (the port numbers are also listed): **afs (1483), bgp (179), biff (512), bootpc (68), bootps (67),** **cmd (514), cvspserver (2401),** **dhcp (67), domain (53),** **eklogin (2105), ekshell (2106), exec (512),** **finger (79), ftp (21), ftp-data (20),** **http (80), https (443),** **ident (113), imap (143),** **kerberos-sec (88), klogin (543), kpasswd (761), krb-prop (754), krbupdate (760), kshell (544),** **ldap (389), login (513),** **mobileip-agent (434), mobilip-mn (435), msdp (639),** **netbios-dgm (138), netbios-ns (137), netbios-ssn (139), nfsd (2049), nntp (119), ntalk (518), ntp (123),** **pop3 (110), pptp (1723), printer (515),** **radacct (1813),radius (1812), rip (520), rkinit (2108),** **smtp (25), snmp (161), snmptrap (162), snpp (444), socks (1080), ssh (22), sunrpc (111), syslog (514),**	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—ports, VLANs, and Layer 3 interfaces	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—ports, VLANs, and Layer 3 interfaces
	tacacs-ds (65), talk (517), telnet (23), tftp (69), timed (525), **who (513),** **xdmcp (177),** **zephyr-clt (2103), zephyr-hm (2104)**		

Table 335: Supported Match Conditions Applicable to IPv4 Traffic for Firewall Filters on J-EX Series Switches *(continued)*

Match Condition	Description	Supported Platforms and Bind Points	
		Ingress	Egress
destination-prefix-list *prefix-list*	IP destination prefix list field. You can define a list of IP address prefixes under a prefix-list alias for frequent use. You make this definition at the **[edit policy-options]** hierarchy level.	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—ports, VLANs, and Layer 3 interfaces	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—ports, VLANs, and Layer 3 interfaces
dot1q-tag *number*	The tag field in the Ethernet header. The tag values can be 1–4095.	• J-EX4200—ports and VLANs • J-EX8200—ports and VLANs	• J-EX4200—ports and VLANs • J-EX8200—not supported
dot1q-user-priority *number*	User-priority field of the tagged Ethernet packet. User-priority values can be 0–7. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): • **background (1)**—Background • **best-effort (0)**—Best effort • **controlled-load (4)**—Controlled load • **excellent-load (3)**—Excellent load • **network-control (7)**—Network control reserved traffic • **standard (2)**—Standard or Spare • **video (5)**—Video • **voice (6)**—Voice	• J-EX4200—ports and VLANs • J-EX8200—ports and VLANs	• J-EX4200—ports and VLANs • J-EX8200—ports and VLANs

Table 335: Supported Match Conditions Applicable to IPv4 Traffic for Firewall Filters on J-EX Series Switches *(continued)*

Match Condition	Description	Supported Platforms and Bind Points	
		Ingress	Egress
dscp *number*	Differentiated Services code point (DSCP). The DiffServ protocol uses the type-of-service (ToS) byte in the IP header. The most significant six bits of this byte form the DSCP. You can specify DSCP in hexadecimal, binary, or decimal form. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): • **ef (46)**—as defined in RFC 2598, *An Expedited Forwarding PHB*. • **af11 (10)**, **af12 (12)**, **af13 (14)**; **af21 (18)**, **af22 (20)**, **af23 (22)**; **af31 (26)**, **af32 (28)**, **af33 (30)**; **af41 (34)**, **af42 (36)**, **af43 (38)** These four classes, with three drop precedences in each class, for a total of 12 code points, are defined in RFC 2597, *Assured Forwarding PHB*.	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—ports, VLANs, and Layer 3 interfaces	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—ports, VLANs, and Layer 3 interfaces
ether-type [aarp | appletalk | arp | ipv4 | ipv6 | mpls—multicast | mpls-unicast | oam | ppp | pppoe-discovery | pppoe-session | sna |*value***]**	Ethernet type field of a packet. The *EtherType value* specifies what protocol is being transported in the Ethernet frame. In place of the numeric value, you can specify one of the following text synonyms: • **aarp**—EtherType value AARP (0x80F3) • **appletalk**—EtherType value AppleTalk (0x809B) • **arp**—EtherType value ARP (0x0806) • **ipv4**—EtherType value IPv4 (0x0800) • **ipv6**—EtherType value IPv6 (0x08DD) • **mpls multicast**—EtherType value MPLS multicast (0x8848) • **mpls unicast**—EtherType value MPLS unicast (0x8847) • **oam**—EtherType value OAM (0x88A8) • **ppp**—EtherType value PPP (0x880B) • **pppoe-discovery**—EtherType value PPPoE Discovery Stage (0x8863) • **pppoe-session**—EtherType value PPPoE Session Stage (0x8864) • **sna**—EtherType value SNA (0x80D5)	• J-EX4200—ports and VLANs • J-EX8200—ports and VLANs	• J-EX4200—ports and VLANs • J-EX8200—not supported.

Table 335: Supported Match Conditions Applicable to IPv4 Traffic for Firewall Filters on J-EX Series Switches *(continued)*

Match Condition	Description	Supported Platforms and Bind Points	
		Ingress	Egress
fragment-flags *fragment-flags*	IP fragmentation flags, specified in symbolic or hexadecimal formats. You can specify one of the following options: dont-fragment (0x4000), more-fragments (0x2000), or reserved (0x8000)	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—ports, VLANs, and Layer 3 interfaces	• J-EX4200—not supported • J-EX8200—not supported
icmp-code *number*	ICMP code field. This value or option provides more specific information than **icmp-type**. Because the value's meaning depends upon the associated **icmp-type**, you must specify **icmp-type** along with **icmp-code**. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed). The options are grouped by the ICMP type with which they are associated: • **parameter-problem—ip-header-bad (0), required-option-missing (1)** • **redirect—redirect-for-host (1), redirect-for-network (0), redirect-for-tos-and-host (3), redirect-for-tos-and-net (2)** • **time-exceeded—ttl-eq-zero-during-reassembly (1), ttl-eq-zero-during-transit (0)** • **unreachable—communication-prohibited-by-filtering (13), destination-host-prohibited (10), destination-host-unknown (7), destination-network-prohibited (9), destination-network-unknown (6), fragmentation-needed (4), host-precedence-violation (14), host-unreachable (1), host-unreachable-for-TOS (12), network-unreachable (0), network-unreachable-for-TOS (11), port-unreachable (3), precedence-cutoff-in-effect (15), protocol-unreachable (2), source-host-isolated (8), source-route-failed (5)**	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—ports, VLANs, and Layer 3 interfaces	• J-EX4200—VLANs and Layer 3 interfaces • J-EX8200—ports, VLANs, and Layer 3 interfaces

Table 335: Supported Match Conditions Applicable to IPv4 Traffic for Firewall Filters on J-EX Series Switches *(continued)*

Match Condition	Description	Supported Platforms and Bind Points	
		Ingress	Egress
icmp-type *number*	ICMP packet type field. Typically, you specify this match in conjunction with the **protocol** match statement to determine which protocol is being used on the port. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): **echo-reply (0)**, **echo-request (8)**, **info-reply (16)**, **info-request (15)**, **mask-request (17)**, **mask-reply (18)**, **parameter-problem (12)**, **redirect (5)**, **router-advertisement (9)**, **router-solicit (10)**, **source-quench (4)**, **time-exceeded (11)**, **timestamp (13)**, **timestamp-reply (14)**, **unreachable (3)**	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—ports, VLANs, and Layer 3 interfaces	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—ports, VLANs, and Layer 3 interfaces
interface *interface-name*	Interface on which the packet is received. You can specify the wildcard character (*) as part of an interface name. NOTE: An interface from which a packet is sent cannot be used as a match condition.	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—ports, VLANs, and Layer 3 interfaces	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—ports, VLANs, and Layer 3 interfaces
ip-options	Presence of the options field in the IP header.	• J-EX4200—Layer 3 interfaces • J-EX8200—Layer 3 interfaces	• J-EX4200—not supported • J-EX8200—not supported
is-fragment	If the packet is a trailing fragment. This match condition does not match the first fragment of a fragmented packet. Use two terms to match both first and trailing fragments.	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—ports, VLANs, and Layer 3 interfaces	• J-EX4200—not supported • J-EX8200—not supported
next-header *bytes*	8-bit protocol field that identifies the type of header immediately following the IPv6 header. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): **ah (51)**, **dstops (60)**, **egp (8)**, **esp (50)**, **fragment (44)**, **gre (47)**, **hop-by-hop (0)**, **icmp (1)**, **icmpv6 (1)**, **igmp (2)**, **ipip (4)**, **ipv6 (41)**, **no-next-header (59)**, **ospf (89)**, **pim (103)**, **routing (43)**, **rsvp (46)**, **sctp (132)**, **tcp (6)**, **udp (17)**, or **vrrp (112)**.	• J-EX4200—Layer 3 interfaces • J-EX8200—not supported	• J-EX4200—Layer 3 interfaces • J-EX8200—not supported

Table 335: Supported Match Conditions Applicable to IPv4 Traffic for Firewall Filters on J-EX Series Switches *(continued)*

Match Condition	Description	Supported Platforms and Bind Points	
		Ingress	Egress
precedence *precedence*	IP precedence. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): **critical-ecp (5)**, **flash (3)**, **flash-override (4)**, **immediate (2)**, **internet-control (6)**, **net-control (7)**, **priority (1)**, or **routine (0)**.	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—ports, VLANs, and Layer 3 interfaces	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—ports, VLANs, and Layer 3 interfaces
protocol *list of protocols*	IPv4 protocol value. In place of the numeric value, you can specify one of the following text synonyms: **egp (8)**, **esp (50)**, **gre (47)**, **icmp (1)**, **igmp (2)**, **ipip (4)**, **ospf (89)**, **pim (103)**, **rsvp (46)**, **tcp (6)**, **udp (17)**	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—ports, VLANs, and Layer 3 interfaces	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—ports, VLANs, and Layer 3 interfaces
source-address *ip-address*	IP source address field, which is the address of the source node sending the packet. For IPV6, the source-address field is 128 bits in length. The filter description syntax supports the text representations for IPv6 addresses that are described in RFC 2373, *IP Version 6 Addressing Architecture*.	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—ports, VLANs, and Layer 3 interfaces	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—ports, VLANs, and Layer 3 interfaces
source-mac-address *mac-address*	Source MAC address.	• J-EX4200—ports and VLANs • J-EX8200—ports and VLANs	• J-EX4200—ports and VLANs • J-EX8200—ports and VLANs
source-port *number*	TCP or UDP **source-port** field. Typically, you specify this match in conjunction with the **protocol** match statement to determine which protocol is being used on the port. In place of the numeric field, you can specify one of the text synonyms listed under **destination-port**.	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—ports, VLANs, and Layer 3 interfaces	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—ports, VLANs, and Layer 3 interfaces
source-prefix-list *prefix-list*	IP source prefix list field. You can define a list of IP address prefixes under a prefix-list alias for frequent use. You make this definition at the **[edit policy-options]** hierarchy level.	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—ports, VLANs, and Layer 3 interfaces	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—ports, VLANs, and Layer 3 interfaces

Table 335: Supported Match Conditions Applicable to IPv4 Traffic for Firewall Filters on J-EX Series Switches *(continued)*

Match Condition	Description	Supported Platforms and Bind Points	
		Ingress	Egress
tcp-established	TCP packets of an established TCP connection. This condition matches packets other than the first packet of a connection. **tcp-established** is a synonym for the bit names **"(ack | rst)"**. **tcp-established** does not implicitly check whether the protocol is TCP. To do so, specify the **protocol tcp** match condition.	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—ports, VLANs, and Layer 3 interfaces	• J-EX4200—not supported • J-EX8200—not supported
tcp-flags [*flags* tcp-initial]	One or more TCP flags: • bit-name—**fin**, **syn**, **rst**, **push**, **ack**, **urgent** • logical operators—**&** (logical AND), **|** (logical OR), **!** (negation) • numerical value—0x01 through 0x20 • text synonym—**tcp-initial** To specify multiple flags, use logical operators.	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—ports, VLANs, and Layer 3 interfaces	• J-EX4200—not supported • J-EX8200—not supported
tcp-initial	Match the first TCP packet of a connection. **tcp-initial** is a synonym for the bit names **"(syn & !ack)"**. **tcp-initial** does not implicitly check whether the protocol is TCP. To do so, specify the **protocol tcp** match condition.	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—ports, VLANs, and Layer 3 interfaces	• J-EX4200—not supported • J-EX8200—not supported

Table 335: Supported Match Conditions Applicable to IPv4 Traffic for Firewall Filters on J-EX Series Switches *(continued)*

Match Condition	Description	Supported Platforms and Bind Points	
		Ingress	Egress
traffic-class	Differentiated Services code point (DSCP). The DiffServ protocol uses the type-of-service (ToS) byte in the IP header. The most significant six bits of this byte form the DSCP. You can specify DSCP in hexadecimal, binary, or decimal form. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): • **ef (46)**—as defined in RFC 2598, *An Expedited Forwarding PHB*. • **af11 (10)**, **af12 (12)**, **af13 (14)**; **af21 (18)**, **af22 (20)**, **af23 (22)**; **af31 (26)**, **af32 (28)**, **af33 (30)**; **af41 (34)**, **af42 (36)**, **af43 (38)** These four classes, with three drop precedences in each class, for a total of 12 code points, are defined in RFC 2597, *Assured Forwarding PHB*.	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—not supported	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—not supported
ttl *value*	TTL type to match. The value can be 1–255.	• J-EX4200—Layer 3 interfaces • J-EX8200—Layer 3 interfaces	• J-EX4200—not supported • J-EX8200—not supported
vlan [*vlan-name* **|** *vlan-id***]**	The VLAN that is associated with the packet.	• J-EX4200—ports and VLANs • J-EX8200—ports and VLANs	• J-EX4200—ports and VLANs • J-EX8200—ports and VLANs

Some of the numeric range and bit-field match conditions allow you to specify a text synonym. For a list of all the synonyms for a match condition, do any of the following:

- If you are using the J-Web Filters Configuration page, select the synonym from the appropriate list.
- If you are using the CLI, type a question mark (**?**) after the **from** statement.

To specify the bit-field value to match, you must enclose the values in quotation marks (**" "**). For example, a match occurs if the RST bit in the TCP flags field is set:

```
tcp-flags "rst";
```

For information about logical operators and how to use bit-field logical operations to create expressions that are evaluated for matches, see “Understanding Firewall Filter Match Conditions” on page 2748.

On J-EX Series Ethernet switches, you can apply a router firewall filter to both IPv4 and IPv6 traffic. You can apply firewall filter match conditions to IPv6 traffic on Layer 3 interfaces, aggregated Ethernet interfaces, and loopback interfaces. Table 336 on page 2738 describes the match conditions you can specify when configuring a firewall filter for IPv6 traffic.

Table 336: Supported Match Conditions Applicable to IPv6 Traffic for Firewall Filters on EX Series Switches

Match Condition	Description	Supported Platforms and Bind Points	
		Ingress	Egress
destination-address *ip-address*	Specifies the 128-bit address that is the final destination node address for the packet. The filter description syntax supports the text representations for IPv6 addresses as described in RFC 2373, *IP Version6 Addressing Architecture*.	• J-EX4200— Layer 3 interfaces • J-EX8200—Layer 3 interfaces	• J-EX4200—Layer 3 interfaces • J-EX8200—Layer 3 interfaces
destination-mac-address *mac-address*	Destination media access control (MAC) address of the packet. You can define a destination MAC address with a prefix, such as **from destination-mac-address 00:01:02:03:04:05/24**. If no prefix is specified, the default value 48 is used.	• J-EX4200—ports and VLANs • J-EX8200—ports and VLANs	• J-EX4200—ports and VLANs • J-EX8200—ports and VLANs

Table 336: Supported Match Conditions Applicable to IPv6 Traffic for Firewall Filters on EX Series Switches *(continued)*

Match Condition	Description	Supported Platforms and Bind Points	
		Ingress	Egress
destination-port *number*	Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) destination port field. Typically, you specify this match in conjunction with the **protocol** match statement to determine which protocol is used on the port. In place of the numeric value, you can specify one of the following text synonyms (the port numbers are also listed): **afs (1483)**, **bgp (179)**, **biff (512)**, **bootpc (68)**, **bootps (67)**, **cmd (514)**, **cvspserver (2401)**, **dhcp (67)**, **domain (53)**, **eklogin (2105)**, **ekshell (2106)**, **exec (512)**, **finger (79)**, **ftp (21)**, **ftp-data (20)**, **http (80)**, **https (443)**, **ident (113)**, **imap (143)**, **kerberos-sec (88)**, **klogin (543)**, **kpasswd (761)**, **krb-prop (754)**, **krbupdate (760)**, **kshell (544)**, **ldap (389)**, **login (513)**, **mobileip-agent (434)**, **mobilip-mn (435)**, **msdp (639)**, **netbios-dgm (138)**, **netbios-ns (137)**, **netbios-ssn (139)**, **nfsd (2049)**, **nntp (119)**, **ntalk (518)**, **ntp (123)**, **pop3 (110)**, **pptp (1723)**, **printer (515)**, **radacct (1813)**,**radius (1812)**, **rip (520)**, **rkinit (2108)**, **smtp (25)**, **snmp (161)**, **snmptrap (162)**, **snpp (444)**, **socks (1080)**, **ssh (22)**, **sunrpc (111)**, **syslog (514)**,	• J-EX4200— VLANs, and Layer 3 interfaces • J-EX8200—Layer 3 interfaces	• J-EX4200—Layer 3 interfaces • J-EX8200—Layer 3 interfaces

Table 336: Supported Match Conditions Applicable to IPv6 Traffic for Firewall Filters on EX Series Switches *(continued)*

Match Condition	Description	Supported Platforms and Bind Points	
		Ingress	Egress
	tacacs-ds (65), **talk (517)**, **telnet (23)**, **tftp (69)**, **timed (525)**, **who (513)**, **xdmcp (177)**, **zephyr-clt (2103)**, **zephyr-hm (2104)**		
destination-prefix-list *prefix-list*	IP destination prefix list field. You can define a list of IP address prefixes under a prefix-list alias for frequent use. You make this definition at the **[edit policy-options]** hierarchy level.	• J-EX4200—Layer 3 interfaces • J-EX8200—Layer 3 interfaces	• J-EX4200—Layer 3 interfaces • J-EX8200—Layer 3 interfaces
dot1q-tag *number*	The tag field in the Ethernet header. The tag values can be 1–4095.	• J-EX4200—ports and VLANs • J-EX8200—ports and VLANs	• J-EX4200—ports and VLANs • J-EX8200—not supported
dot1q-user-priority *number*	User-priority field of the tagged Ethernet packet. User-priority values can be 0–7. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): • **background (1)**—Background • **best-effort (0)**—Best effort • **controlled-load (4)**—Controlled load • **excellent-load (3)**—Excellent load • **network-control (7)**—Network control reserved traffic • **standard (2)**—Standard or Spare • **video (5)**—Video • **voice (6)**—Voice	• J-EX4200—ports and VLANs • J-EX8200—ports and VLANs	• J-EX4200—ports and VLANs • J-EX8200—ports and VLANs
ether-type (ipv6)*value*	Ethernet type field of a packet. The EtherType value specifies what protocol is being transported in the Ethernet frame. In place of the numeric value, you can specify the following text synonym: • **ipv6**—EtherType value IPv6 (0x08DD)	• J-EX4200—ports and VLANs • J-EX8200—ports and VLANs	• J-EX4200—ports and VLANs • J-EX8200— ports and VLANs.

Table 336: Supported Match Conditions Applicable to IPv6 Traffic for Firewall Filters on EX Series Switches *(continued)*

Match Condition	Description	Supported Platforms and Bind Points	
		Ingress	Egress
icmp-code *number*	ICMP code field. This value or option provides more specific information than **icmp-type**. Because the value's meaning depends upon the associated **icmp-type**, you must specify **icmp-type** along with **icmp-code**. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed). The options are grouped by the ICMP type with which they are associated: • **parameter-problem**—**ip-header-bad (0)**, **unrecognized-next-header (1)**, **unrecognized-option (2)** • **time-exceeded**—**ttl-eq-zero-during-reassembly (1)**, **ttl-eq-zero-during-transit (0)** • **destination-unreachable**—**no-route-to--destination (0)**, **administratively-prohibited (1)**, **address-unreachable (3)**, **port-unreachable (4)**	• J-EX4200—Layer 3 interfaces • J-EX8200—Layer 3 interfaces	• J-EX4200—Layer 3 interfaces • J-EX8200—Layer 3 interfaces
icmp-type *number*	ICMP packet type field. Typically, you specify this match in conjunction with the **protocol** match statement to determine which protocol is being used on the port. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): **echo-reply (0)**, **echo-request (8)**, **info-reply (16)**, **info-request (15)**, **mask-request (17)**, **mask-reply (18)**, **parameter-problem (12)**, **redirect (5)**, **router-advertisement (9)**, **router-solicit (10)**, **source-quench (4)**, **time-exceeded (11)**, **timestamp (13)**, **timestamp-reply (14)**, **unreachable (3)**	• J-EX4200—Layer 3 interfaces • J-EX8200—Layer 3 interfaces	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—Layer 3 interfaces
interface *interface-name*	Interface on which the packet is received.	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—ports, VLANs, and Layer 3 interfaces	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—ports, VLANs, and Layer 3 interfaces

Table 336: Supported Match Conditions Applicable to IPv6 Traffic for Firewall Filters on EX Series Switches *(continued)*

Match Condition	Description	Supported Platforms and Bind Points	
		Ingress	Egress
next-header *bytes*	8-bit protocol field that identifies the type of header immediately following the IPv6 header. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): **ah (51)**, **dstops (60)**, **egp (8)**, **esp (50)**, **fragment (44)**, **gre (47)**, **hop-by-hop (0)**, **icmp (1)**, **icmpv6 (1)**, **igmp (2)**, **ipip (4)**, **ipv6 (41)**, **no-next-header (59)**, **ospf (89)**, **pim (103)**, **routing (43)**, **rsvp (46)**, **sctp (132)**, **tcp (6)**, **udp (17)**, or **vrrp (112)**.	• J-EX4200—Layer 3 interfaces • J-EX8200—Layer 3 interfaces	• J-EX4200—Layer 3 interfaces • J-EX8200—Layer 3 interfaces
packet-length *bytes*	Length of the received packet, in bytes. The length refers only to the IP packet, including the packet header, and does not include any Layer 2 encapsulation overhead.	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—Layer 3 interfaces	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—Layer 3 interfaces
source-address *ip-address*	IP source address field, which is 128 bits in length. The filter description syntax supports the text representations for IPv6 addresses that are described in RFC 2373, *IP Version 6 Addressing Architecture*.	• J-EX4200—Layer 3 interfaces • J-EX8200—Layer 3 interfaces	• J-EX4200—Layer 3 interfaces • J-EX8200—Layer 3 interfaces
source-mac-address *mac-address*	Source MAC address. You can define a source MAC address with a prefix, such as **from destination-mac-address 00:01:02:03:04:05/24**. If no prefix is specified, the default value 48 is used.	• J-EX4200—ports and VLANs • J-EX8200—ports and VLANs	• J-EX4200—ports and VLANs • J-EX8200—ports and VLANs
source-port *number*	TCP or UDP **source-port** field. Typically, you specify this match in conjunction with the **next-header** match statement to determine which next-header is being used on the port. In place of the numeric field, you can specify one of the text synonyms listed under **destination-port**.	• J-EX4200—Layer 3 interfaces • J-EX8200—Layer 3 interfaces	• J-EX4200—Layer 3 interfaces • J-EX8200—Layer 3 interfaces
source-prefix-list *prefix-list*	IP source prefix list field. You can define a list of IP address prefixes under a prefix-list alias for frequent use. You make this definition at the **[edit policy-options]** hierarchy level.	• J-EX4200—Layer 3 interfaces • J-EX8200—Layer 3 interfaces	• J-EX4200—Layer 3 interfaces • J-EX8200—Layer 3 interfaces

Table 336: Supported Match Conditions Applicable to IPv6 Traffic for Firewall Filters on EX Series Switches *(continued)*

Match Condition	Description	Supported Platforms and Bind Points	
		Ingress	Egress
tcp-flags (*flags* tcp-initial)	One or more TCP flags: • bit-name—**fin**, **syn**, **rst**, **push**, **ack**, **urgent** • logical operators—**&** (logical AND), **|** (logical OR), **!** (negation) • numerical value—0x01 through 0x20 • text synonym—**tcp-initial** To specify multiple flags, use logical operators.	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—Layer 3 interfaces	• J-EX4200—not supported • J-EX8200—not supported
tcp-initial	Match the first TCP packet of a connection. **tcp-initial** is a synonym for the bit names **"(syn & !ack)"**. **tcp-initial** does not implicitly check whether the protocol is TCP. To do so, specify the **protocol tcp** match condition.	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—Layer 3 interfaces	• J-EX4200—not supported • J-EX8200—not supported
traffic-class *number*	Differentiated Services code point (DSCP). The DiffServ protocol uses the type-of-service (ToS) byte in the IP header. The most significant six bits of this byte form the DSCP. You can specify DSCP in hexadecimal, binary, or decimal form. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed): • **ef (46)**—as defined in RFC 2598, *An Expedited Forwarding PHB*. • **af11 (10)**, **af12 (12)**, **af13 (14)**; **af21 (18)**, **af22 (20)**, **af23 (22)**; **af31 (26)**, **af32 (28)**, **af33 (30)**; **af41 (34)**, **af42 (36)**, **af43 (38)** These four classes, with three drop precedences in each class, for a total of 12 code points, are defined in RFC 2597, *Assured Forwarding PHB*.	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—Layer 3 interfaces	• J-EX4200—ports, VLANs, and Layer 3 interfaces • J-EX8200—Layer 3 interfaces
vlan (*vlan-id* | *vlan-name*)	The VLAN that is associated with the packet.	• J-EX4200—ports and VLANs • J-EX8200—ports and VLANs	• J-EX4200—ports and VLANs • J-EX8200—ports and VLANs

When you define one or more terms that specify the filtering criteria, you also define the action to take if the packet matches all criteria. Table 337 on page 2744 shows the actions that you can specify in a term.

Table 337: Actions for Firewall Filters

Action	Description	Supported Platforms and Direction
accept	Accept a packet.	• J-EX4200—ingress and egress • J-EX8200—ingress and egress
discard	Discard a packet silently without sending an Internet Control Message Protocol (ICMP) message.	• J-EX4200—ingress and egress • J-EX8200—ingress and egress
reject *message-type*	Discard a packet, and send an ICMPv4 message (type 3) "destination unreachable". You can log the rejected packets if you configure the **syslog** action modifier. You can specify one of the following message codes: **administratively-prohibited (default), bad-host-tos, bad-network-tos, host-prohibited, host-unknown, host-unreachable, network-prohibited, network-unknown, network-unreachable, port-unreachable, precedence-cutoff, precedence-violation, protocol-unreachable, source-host-isolated, source-route-failed,** or **tcp-reset**. If you specify **tcp-reset**, a TCP reset is returned if the packet is a TCP packet. Otherwise nothing is returned. If you do not specify a message type, the ICMP notification "destination unreachable" is sent with the default message "communication administratively filtered". NOTE: **reject** is not a supported action for IPv6 traffic.	• J-EX4200—ingress only • J-EX8200—ingress only
routing-instance *routing-instance-name*	Forward matched packets to a virtual routing instance.	• J-EX4200—ingress and egress • J-EX8200—not supported
vlan *vlan-name*	Forward matched packets to a specific VLAN. NOTE: **vlan** is not a supported action for IPv6 traffic.	• J-EX4200—ingress only • J-EX8200—not supported

In addition to the actions, you can specify action modifiers. Table 338 on page 2744 shows the action modifiers that you can specify in a term.

Table 338: Action Modifiers for Firewall Filters

Action Modifier	Description	Supported Platforms and Direction
analyzer *analyzer-name*	Mirror port traffic to a specified destination port or VLAN that is connected to a protocol analyzer application. Mirroring copies all packets seen on one switch port to a network monitoring connection on another switch port. The analyzer name must be configured under **[edit ethernet-switching-options analyzer]**.	• J-EX4200—ingress only • J-EX8200—ingress only

Table 338: Action Modifiers for Firewall Filters *(continued)*

Action Modifier	Description	Supported Platforms and Direction
count *counter-name*	Count the number of packets that pass this filter, term, or policer.	• J-EX4200—ingress and egress • J-EX8200—not supported
forwarding-class *class*	Classify the packet in one of the following forwarding classes: • **assured-forwarding** • **best-effort** • **expedited-forwarding** • **network-control**	• J-EX4200—ingress and egress • J-EX8200—ingress and egress
interface *interface-name*	Forward the traffic to the specified interface bypassing the switching lookup. NOTE: **interface** is not a supported action modifier for IPv6 traffic.	• J-EX4200—ingress only • J-EX8200—ingress only
log	Log the packet's header information in the Routing Engine. To view this information, issue the **show firewall log** command in the CLI. NOTE: **log** is not a supported action modifier for IPv6 traffic.	• J-EX4200—ingress only • J-EX8200—ingress only
loss-priority (*high* | *low*)	Set the packet loss priority (PLP).	• J-EX4200—ingress only • J-EX8200—not supported
policer *policer-name*	Apply rate limits to the traffic. You can specify a policer for ingress port, VLAN, and router firewall filters only.	• J-EX4200—ingress only • J-EX8200—ingress only
syslog	Log an alert for this packet. You can specify that the log be sent to a server for storage and analysis. NOTE: **syslog** is not a supported action modifier for IPv6 traffic.	• J-EX4200—ingress only • J-EX8200—ingress only

NOTE: On J-EX Series switches, accept and discard are the only actions supported for firewall filters applied on loopback interfaces.

Related Documentation

- Firewall Filter Configuration Statements Supported by Junos OS for J-EX Series Switches on page 2806
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on J-EX Series Switches on page 2773

- Understanding Firewall Filter Match Conditions on page 2748
- Understanding How Firewall Filters Are Evaluated on page 2746
- Understanding How Firewall Filters Test a Packet's Protocol on page 2752
- Understanding the Use of Policers in Firewall Filters on page 2752
- Understanding Filter-Based Forwarding for J-EX Series Switches on page 2753

Understanding How Firewall Filters Are Evaluated

A firewall filter consists of one or more terms, and the order of the terms within a firewall filter is important. Before you configure firewall filters, you should understand how J-EX Series Switches evaluate the terms within a firewall filter and how packets are evaluated against the terms.

When a firewall filter consists of a single term, the filter is evaluated as follows:

- If the packet matches all the conditions, the action in the **then** statement is taken.
- If the packet matches all the conditions, and no action is specified in the **then** statement, the default action **accept** is taken.

When a firewall filter consists of more than one term, the firewall filter is evaluated sequentially:

1. The packet is evaluated against the conditions in the **from** statement in the first term.
2. If the packet matches all the conditions in the term, the action in the **then** statement is taken and the evaluation ends. Subsequent terms in the filter are not evaluated.
3. If the packet does not match all the conditions in the term, the packet is evaluated against the conditions in the **from** statement in the second term.

   This process continues until either the packet matches the conditions in the **from** statement in one of the subsequent terms or there are no more terms in the filter.
4. If a packet passes through all the terms in the filter without a match, the packet is discarded.

Figure 77 on page 2747 shows how a J-EX Series switch evaluates the terms within a firewall filter.

Figure 77: Evaluation of Terms Within a Firewall Filter

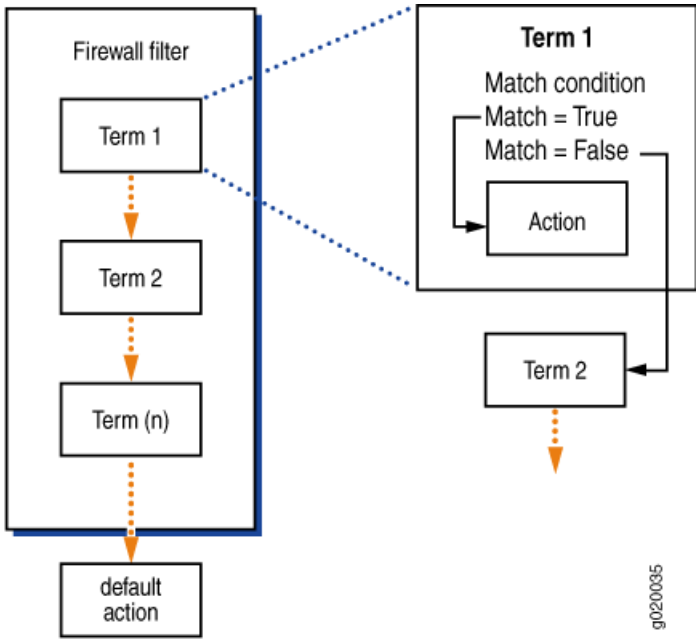

If a term does not contain a **from** statement, the packet is considered to match and the action in the **then** statement of the term is taken.

If a term does not contain a **then** statement, or if an action has not been configured in the **then** statement, and the packet matches the conditions in the **from** statement of the term, the packet is accepted.

Every firewall filter contains an implicit **deny** statement at the end of the filter, which is equivalent to the following explicit filter term:

```
term implicit-rule {
then discard;
}
```

Consequently, if a packet passes through all the terms in a filter without matching any conditions, the packet is discarded. If you configure a firewall filter that has no terms, all packets that pass through the filter are discarded.

NOTE: Firewall filtering is supported on packets that are at least 48 bytes long.

Related Documentation

- Firewall Filters for J-EX Series Switches Overview on page 2721
- Understanding Firewall Filter Match Conditions on page 2748
- Understanding the Use of Policers in Firewall Filters on page 2752
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755

Understanding Firewall Filter Match Conditions

Before you define terms for firewall filters, you must understand how the conditions that you specify in a term are handled and how to specify interface filter, numeric filter, address filter, and bit-field filter match conditions to achieve the desired filtering results.

Filter Match Conditions

In the **from** statement of a firewall filter term, you specify the conditions that the packet must match for the action in the **then** statement to be taken. All conditions in the **from** statement must match for the action to be taken. The order in which you specify match conditions is not important, because a packet must match all the conditions in a term for a match to occur.

If you specify no match conditions in a term, that term matches all packets.

An individual condition in a **from** statement cannot contain a list of values. For example, you cannot specify numeric ranges or multiple source or destination addresses.

Individual conditions in a **from** statement cannot be negated. A negated condition is an explicit mismatch.

Numeric Filter Match Conditions

Numeric filter conditions match packet fields that are identified by a numeric value, such as port and protocol numbers. For numeric filter match conditions, you specify a keyword that identifies the condition and a single value that a field in a packet must match.

You can specify the numeric value in one of the following ways:

- Single number—A match occurs if the value of the field matches the number. For example:

  source-port 25;

- Text synonym for a single number— A match occurs if the value of the field matches the number that corresponds to the synonym. For example:

  source-port http;

To specify more than one value in a filter term, you enter each value in its own match statement. For example, a match occurs in the following term if the value of **vlan** field is 10 or 30.

[edit firewall family *family-name* filter *filter-name* term *term-name* from]

```
vlan 10;
vlan 30;
```

The following restrictions apply to numeric filter match conditions:

- You cannot specify a range of values.
- You cannot specify a list of comma-separated values.
- You cannot exclude a specific value in a numeric filter match condition. For example, you cannot specify a condition that would match only if the match condition was not equal to a given value.

Interface Filter Match Conditions

Interface filter match conditions can match interface name values in a packet. For interface filter match conditions, you specify the name of the interface, for example:

```
[edit firewall family family-name filter filter-name term term-name from]
user@host# set interface ge-0/0/1
```

Port and VLAN interfaces do not use logical unit numbers. However, a firewall filter that is applied to a router interface can specify the logical unit number in the interface filter match condition, for example:

```
[edit firewall family family-name filter filter-name term term-name from]
user@host# set interface ge-0/1/0.0
```

You can include the * wildcard as part of the interface name, for example:

```
[edit firewall family family-name filter filter-name term term-name from]
user@host# set interface ge-0/*/1
user@host# set interface ge-0/1/*
user@host# set interface ge-*
```

IP Address Filter Match Conditions

Address filter match conditions can match prefix values in a packet, such as IP source and destination prefixes. For address filter match conditions, you specify a keyword that identifies the field and one prefix of that type that a packet must match.

You specify the address as a single prefix. A match occurs if the value of the field matches the prefix. For example:

```
[edit firewall family family-name filter filter-name term term-name from]
user@host# set destination-address 10.2.1.0/28;
```

Each prefix contains an implicit 0/0 except statement, which means that any prefix that does not match the prefix that is specified is explicitly considered not to match.

To specify the address prefix, use the notation prefix/prefix-length. If you omit prefix-length, it defaults to /32. For example:

```
[edit firewall family family-name filter filter-name term term-name from]
user@host# set destination-address 10
[edit firewall family family-name filter filter-name term term-name from]
user@host# show
destination-address {
10.0.0.0/32;
}
```

To specify more than one IP address in a filter term, you enter each address in its own match statement. For example, a match occurs in the following term if the value of the **source-address** field matches either of the following source-address prefixes:

```
[edit firewall family family-name filter filter-name term term-name from]
user@host# set source-address 10.0.0.0/8
user@host# set source-address 10.1.0.0/16
```

MAC Address Filter Match Conditions

MAC address filter match conditions can match source and destination MAC address values in a packet. For MAC address filter match conditions, you specify a keyword that identifies the field and one value of that type that a packet must match.

You can specify the MAC address as six hexadecimal bytes in the following formats:

```
[edit firewall family family-name filter filter-name term term-name from]
user@host# set destination-mac-address 0011.2233.4455
```

```
[edit firewall family family-name filter filter-name term term-name from]
user@host# set destination-mac-address 00:11:22:33:44:55
```

```
[edit firewall family family-name filter filter-name term term-name from]
user@host# set destination-mac-address 001122334455
```

To specify more than one MAC address in a filter term, you enter each MAC address in its own match statement. For example, a match occurs in the following term if the value of the **source-mac-address** field matches either of the following addresses.

```
[edit firewall family family-name filter filter-name term term-name from]
user@host# set source-mac-address 00:11:22:33:44:55
user@host# set source-mac-address 00:11:22:33:20:15
```

Bit-Field Filter Match Conditions

Bit-field filter conditions match packet fields if particular bits in those fields are or are not set. You can match the IP options, TCP flags, and IP fragmentation fields. For bit-field filter match conditions, you specify a keyword that identifies the field and tests to determine that the option is present in the field.

To specify the bit-field value to match, enclose the value in double quotation marks. For example, a match occurs if the **RST** bit in the TCP flags field is set:

```
[edit firewall family family-name filter filter-name term term-name from]
user@host# set tcp-flags "rst"
```

Typically, you specify the bits to be tested by using keywords. Bit-field match keywords always map to a single bit value. You also can specify bit fields as hexadecimal or decimal numbers.

To match multiple bit-field values, use the logical operators, which are described in Table 339 on page 2750. The operators are listed in order from highest precedence to lowest precedence. Operations are left-associative.

Table 339: Actions for Firewall Filters

Logical Operators	Description
!	Negation.

Table 339: Actions for Firewall Filters *(continued)*

Logical Operators	Description
& or +	Logical AND.
| or ,	Logical OR.

To negate a match, precede the value with an exclamation point. For example, a match occurs only if the RST bit in the TCP flags field is not set:

```
[edit firewall family family-name filter filter-name term term-name from]
user@host# set tcp-flags "!rst"
```

In the following example of a logical AND operation, a match occurs if the packet is the initial packet on a TCP session:

```
[edit firewall family family-name filter filter-name term term-name from]
user@host# set tcp-flags "syn&!ack"
```

In the following example of a logical OR operation, a match occurs if the packet is not the initial packet on a TCP session:

```
[edit firewall family family-name filter filter-name term term-name from]
user@host# set tcp-flags "syn|ack"
```

For a logical OR operation, you can specify a maximum of two match conditions in a single term. If you need to match more than two bit-field values in a logical OR operation, configure the same match condition in consecutive terms with additional bit-field values. In the following example, the two terms configured match the SYN, ACK, FIN, or RST bit in the TCP flags field:

```
[edit firewall family family-name filter filter-name term term-name1 from]
user@host# set tcp-flags "syn|ack"
[edit firewall family family-name filter filter-name term term-name2 from]
user@host# set tcp-flags "fin|rst"
```

You can use text synonyms to specify some common bit-field matches. You specify these matches as a single keyword. In the following example of a text synonym, a match occurs if the packet is the initial packet on a TCP session:

```
[edit firewall family family-name filter filter-name term term-name from]
user@host# set tcp-flags tcp-initial
```

Related Documentation

- Firewall Filters for J-EX Series Switches Overview on page 2721
- Understanding How Firewall Filters Test a Packet's Protocol on page 2752
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on J-EX Series Switches on page 2773
- Firewall Filter Match Conditions and Actions for J-EX Series Switches on page 2728

Understanding How Firewall Filters Test a Packet's Protocol

When examining match conditions, the Junos OS for J-EX Series Switches tests only the field that is specified. The software does not implicitly test the IP header to determine whether a packet is an IP packet. Therefore, in some cases, you must specify **protocol** field match conditions in conjunction with other match conditions to ensure that the filters are performing the expected matches.

If you specify a protocol match condition or a match of the ICMP type or TCP flags field, there is no implied protocol match. For the following match conditions, you must explicitly specify the protocol match condition in the same term:

- **destination-port**—Specify the match **protocol tcp** or **protocol udp**.
- **source-port**—Specify the match **protocol tcp** or **protocol udp**.

If you do not specify the protocol when using the preceding fields, design your filters carefully to ensure that they perform the expected matches. For example, if you specify a match of **destination-port ssh**, the switch deterministically matches any packets that have a value of **22** in the two-byte field that is two bytes beyond the end of the IP header without ever checking the IP protocol field.

Related Documentation

- Firewall Filters for J-EX Series Switches Overview on page 2721
- Understanding Firewall Filter Match Conditions on page 2748
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755

Understanding the Use of Policers in Firewall Filters

Policing, or rate limiting, is an important component of firewall filters that lets you control the amount of traffic that enters an interface.

A single firewall filter configured with a policer permits only traffic at specified data rates to provide protection from denial-of-service (DOS) attacks. Traffic that exceeds the rate limits specified by the policer can be discarded. Discard is the only supported policer action. Typically, traffic that exceeds the rate limits specified by the policer is either discarded or marked as lower priority than traffic that meets the rate limits specified by the policer. When necessary, low-priority traffic can be discarded by the switch to prevent congestion.

A policer applies two types of rate limits on traffic:

- Bandwidth—The number of bits per second permitted, on average
- Maximum burst size—The maximum size permitted for bursts of data that exceed the given bandwidth limit

Policing uses an algorithm to enforce a limit on average bandwidth while allowing bursts up to a specified maximum value. You can define specific classes of traffic on an interface

and apply a set of rate limits to each class. After you name and configure a policer, it is stored as a template. You can then use a policer in a firewall filter configuration.

Each policer you configure includes an implicit counter that counts the number of packets exceeding the rate limits specified for the policer. To get filter or term-specific packets counts, you must configure a new policer for each filter or term that requires policing.

Related Documentation

- Firewall Filters for J-EX Series Switches Overview on page 2721
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Firewall Filter Match Conditions and Actions for J-EX Series Switches on page 2728

Understanding Filter-Based Forwarding for J-EX Series Switches

Administrators of J-EX Series Switches can use firewall filters in conjunction with virtual routing instances to specify different routes for packets to travel in their networks. To set up this feature, which is called filter-based forwarding, you specify a filter and match criteria and then specify the virtual routing instance to send packets to.

You might want to use filter-based forwarding to route specific types of traffic through a firewall or security device before the traffic continues on its path. You can also use filter-based forwarding to give certain types of traffic preferential treatment or to improve load balancing of switch traffic.

Related Documentation

- Understanding Virtual Routing Instances on J-EX Series Switches on page 1048
- Firewall Filters for J-EX Series Switches Overview on page 2721
- Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on J-EX Series Switches on page 2773

CHAPTER 101

Examples of Firewall Filters Configuration

Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches

This example shows how to configure and apply firewall filters to control traffic that is entering or exiting a port on the switch, a VLAN on the network, and a Layer 3 interface on the switch. Firewall filters define the rules that determine whether to forward or deny packets at specific processing points in the packet flow.

Requirements

This example uses the following software and hardware components:

- Two J-EX4200-48T switches: one to be used as an access switch, the other to be used as a distribution switch
- One uplink module
- One router

Before you configure and apply the firewall filters in this example, be sure you have:

- An understanding of firewall filter concepts, policers, and CoS
- Installed the uplink module in the distribution switch. See Installing an Uplink Module in a J-EX4200 Switch.

Overview

This configuration example show how to configure and apply firewall filters to provide rules to evaluate the contents of packets and determine when to discard, forward, classify, count, and analyze packets that are destined for or originating from the J-EX Series switches that handle all **voice-vlan**, **employee-vlan**, and **guest-vlan** traffic. Table 340 on page 2756 shows the firewall filters that are configured for the J-EX Series switches in this example.

Table 340: Configuration Components: Firewall Filters

Component	Purpose/Description
Port firewall filter, **ingress-port-voip-class-limit-tcp-icmp**	This firewall filter performs two functions: • Assigns priority queueing to packets with a source MAC address that matches the phone MAC addresses. The forwarding class **expedited-forwarding** provides low loss, low delay, low jitter, assured bandwidth, and end-to-end service for all **voice-vlan** traffic. • Performs rate limiting on packets that enter the ports for **employee-vlan**. The traffic rate for TCP and ICMP packets is limited to 1 Mbps with a burst size up to 30,000 bytes. This firewall filter is applied to port interfaces on the access switch.
VLAN firewall filter, **ingress-vlan-rogue-block**	Prevents rogue devices from using HTTP sessions to mimic the gatekeeper device that manages call registration, admission, and call status for VoIP calls. Only TCP or UDP ports should be used; and only the gatekeeper uses HTTP. That is, all **voice-vlan** traffic on TCP ports should be destined for the gatekeeper device. This firewall filter applies to all phones on **voice-vlan**, including communication between any two phones on the VLAN and all communication between the gatekeeper device and VLAN phones. This firewall filter is applied to VLAN interfaces on the access switch.
VLAN firewall filter, **egress-vlan-watch-employee**	Accepts **employee-vlan** traffic destined for the corporate subnet, but does not monitor this traffic. Employee traffic destined for the Web is counted and analyzed. This firewall filter is applied to vlan interfaces on the access switch.
VLAN firewall filter, **ingress-vlan-limit-guest**	Prevents guests (non-employees) from talking with employees or employee hosts on **employee-vlan**. Also prevents guests from using peer-to-peer applications on **guest-vlan**, but allows guests to access the Web. This firewall filter is applied to VLAN interfaces on the access switch.
Router firewall filter, **egress-router-corp-class**	Prioritizes **employee-vlan** traffic, giving highest forwarding-class priority to employee traffic destined for the corporate subnet. This firewall filter is applied to a routed port (Layer 3 uplink module) on the distribution switch.

Figure 78 on page 2757 shows the application of port, VLAN, and Layer 3 routed firewall filters on the switch.

Figure 78: Application of Port, VLAN, and Layer 3 Routed Firewall Filters

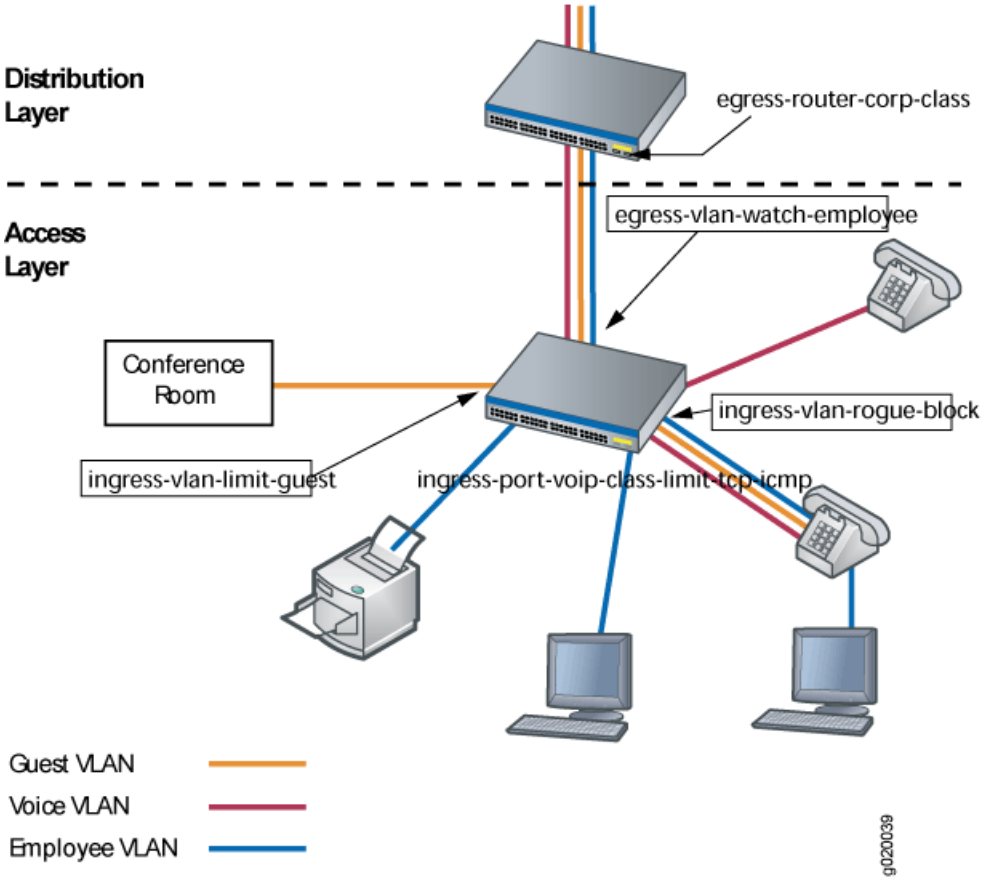

Network Topology

The topology for this configuration example consists of one J-EX4200-48T switch at the access layer, and one J-EX4200-48T switch at the distribution layer. The distribution switch's uplink module is configured to support a Layer 3 connection to a J-series router.

The J-EX Series switches are configured to support VLAN membership. Table 341 on page 2757 shows the VLAN configuration components for the VLANs.

Table 341: Configuration Components: VLANs

VLAN Name	VLAN ID	VLAN Subnet and Available IP Addresses	VLAN Description
voice-vlan	10	192.0.2.0/28 192.0.2.1 through 192.0.2.14 192.0.2.15 is subnet's broadcast address	Voice VLAN used for employee VoIP traffic

Table 341: Configuration Components: VLANs *(continued)*

VLAN Name	VLAN ID	VLAN Subnet and Available IP Addresses	VLAN Description
employee-vlan	**20**	**192.0.2.16/28 192.0.2.17** through **192.0.2.30 192.0.2.31** is subnet's broadcast address	VLAN standalone PCs, PCs connected to the network through the hub in VoIP telephones, wireless access points, and printers. This VLAN completely includes the voice VLAN. Two VLANs (**voice-vlan** and **employee-vlan**) must be configured on the ports that connect to the telephones.
guest-vlan	**30**	**192.0.2.32/28 192.0.2.33** through **192.0.2.46 192.0.2.47** is subnet's broadcast address	VLAN for guests' data devices (PCs). The scenario assumes that the corporation has an area open to visitors, either in the lobby or in a conference room, that has a hub to which visitors can plug in their PCs to connect to the Web and to their company's VPN.
camera-vlan	**40**	**192.0.2.48/28 192.0.2.49** through **192.0.2.62 192.0.2.63** is subnet's broadcast address	VLAN for the corporate security cameras.

Ports on the J-EX Series switches support Power over Ethernet (PoE) to provide both network connectivity and power for VoIP telephones connecting to the ports. Table 342 on page 2758 shows the switch ports that are assigned to the VLANs and the IP and MAC addresses for devices connected to the switch ports:

Table 342: Configuration Components: Switch Ports on a 48-Port All-PoE Switch

Switch and Port Number	VLAN Membership	IP and MAC Addresses	Port Devices
ge-0/0/0, ge-0/0/1	**voice-vlan**, **employee-vlan**	IP addresses: **192.0.2.1** through **192.0.2.2** MAC addresses: **00.05.85.00.00.01**, **00.05.85.00–00.02**	Two VoIP telephones, each connected to one PC.
ge-0/0/2, ge-0/0/3	**employee-vlan**	**192.0.2.17** through **192.0.2.18**	Printer, wireless access points
ge-0/0/4, ge-0/0/5	**guest-vlan**	**192.0.2.34** through **192.0.2.35**	Two hubs into which visitors can plug in their PCs. Hubs are located in an area open to visitors, such as a lobby or conference room
ge-0/0/6, ge-0/0/7	**camera-vlan**	**192.0.2.49** through **192.0.2.50**	Two security cameras

Table 342: Configuration Components: Switch Ports on a 48-Port All-PoE Switch *(continued)*

Switch and Port Number	VLAN Membership	IP and MAC Addresses	Port Devices
ge-0/0/9	**voice-vlan**	IP address: **192.0.2.14** MAC address:**00.05.85.00.00.0E**	Gatekeeper device. The gatekeeper manages call registration, admission, and call status for VoIP phones.
ge-0/1/0		IP address: **192.0.2.65**	Layer 3 connection to a router; note that this is a port on the switch's uplink module

Configuring an Ingress Port Firewall Filter to Prioritize Voice Traffic and Rate-Limit TCP and ICMP Traffic

To configure and apply firewall filters for port, VLAN, and router interfaces, perform these tasks:

CLI Quick Configuration

To quickly configure and apply a port firewall filter to prioritize voice traffic and rate-limit packets that are destined for the **employee-vlan** subnet, copy the following commands and paste them into the switch terminal window:

```
[edit]
set firewall policer tcp-connection-policer if-exceeding burst-size-limit 30k bandwidth-limit 1m
set firewall policer tcp-connection-policer then discard
set firewall policer icmp-connection-policer if-exceeding burst-size-limit 30k bandwidth-limit
1m
set firewall policer icmp-connection-policer then discard
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term voip-high
from source-mac-address 00.05.85.00.00.01
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term voip-high
from source-mac-address 00.05.85.00.00.02
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term voip-high
from protocol udp
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term voip-high
then forwarding-class expedited-forwarding
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term voip-high
then loss-priority low
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
network-control from precedence net-control
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
network-control then forwarding-class network-control
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
network-control then loss-priority low
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
tcp-connection from destination-address 192.0.2.16/28
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
tcp-connection from protocol tcp
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
tcp-connection then policer tcp-connection-policer
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
tcp-connection then count tcp-counter
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
tcp-connection then forwarding-class best-effort
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
tcp-connection then loss-priority high
```

```
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
icmp-connection from destination-address 192.0.2.16/28
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
icmp-connection from protocol icmp
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
icmp-connection then policer icmp-connection-policer
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
icmp-connection then count icmp-counter
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
icmp-connection then forwarding-class best-effort
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term
icmp-connection then loss-priority high
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term best-effort
then forwarding-class best-effort
set firewall family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp term best-effort
then loss-priority high
set interfaces ge-0/0/0 description "voice priority and tcp and icmp traffic rate-limiting filter at
ingress port"
set interfaces ge-0/0/0 unit 0 family ethernet-switching filter input
ingress-port-voip-class-limit-tcp-icmp
set interfaces ge-0/0/1 description "voice priority and tcp and icmp traffic rate-limiting filter at
ingress port"
set interfaces ge-0/0/1 unit 0 family ethernet-switching filter input
ingress-port-voip-class-limit-tcp-icmp
set class-of-service schedulers voice-high buffer-size percent 15
set class-of-service schedulers voice-high priority high
set class-of-service schedulers net-control buffer-size percent 10
set class-of-service schedulers net-control priority high
set class-of-service schedulers best-effort buffer-size percent 75
set class-of-service schedulers best-effort priority low
set class-of-service scheduler-maps ethernet-diffsrv-cos-map forwarding-class
expedited-forwarding scheduler voice-high
set class-of-service scheduler-maps ethernet-diffsrv-cos-map forwarding-class network-control
scheduler net-control
set class-of-service scheduler-maps ethernet-diffsrv-cos-map forwarding-class best-effort
scheduler best-effort
```

Step-by-Step Procedure

To configure and apply a port firewall filter to prioritize voice traffic and rate-limit packets that are destined for the **employee-vlan** subnet:

1. Define the policers **tcp-connection-policer** and **icmp-connection-policer**:

```
[edit]
user@switch# set firewall policer tcp-connection-policer if-exceeding burst-size-limit
30k bandwidth-limit 1m
user@switch# set firewall policer tcp-connection-policer then discard
user@switch# set firewall policer icmp-connection-policer if-exceeding burst-size-limit
30k bandwidth-limit 1m
user@switch# set firewall policer icmp-connection-policer then discard
```

2. Define the firewall filter **ingress-port-voip-class-limit-tcp-icmp**:

```
[edit firewall]
user@switch# set family ethernet-switching filter ingress-port-voip-class-limit-tcp-icmp
```

3. Define the term **voip-high**:

```
[edit firewall family ethernet-switching filter
ingress-port-voip-class-limit-tcp-icmp ]
user@switch# set term voip-high from source-mac-address 00.05.85.00.00.01
user@switch# set term voip-high from source-mac-address 00.05.85.00.00.02
```

```
user@switch# set term voip-high from protocol udp
user@switch# set term voip-high then forwarding-class expedited-forwarding
user@switch# set term voip-high then loss-priority low
```

4. Define the term **network-control**:

```
[edit firewall family ethernet-switching filter
ingress-port-voip-class-limit-tcp-icmp ]
user@switch# set term network-control from precedence net-control
user@switch# set term network-control then forwarding-class network-control
user@switch# set term network-control then loss-priority low
```

5. Define the term **tcp-connection** to configure rate limits for TCP traffic:

```
[edit firewall family ethernet-switching filter
ingress-port-voip-class-limit-tcp-icmp]
user@switch# set term tcp-connection from destination-address 192.0.2.16/28
user@switch# set term tcp-connection from protocol tcp
user@switch# set term tcp-connection then policer tcp-connection-policer
user@switch# set term tcp-connection then count tcp-counter
user@switch# set term tcp-connection then forwarding-class best-effort
user@switch# set term tcp-connection then loss-priority high
```

6. Define the term **icmp-connection** to configure rate limits for ICMP traffic:

```
[edit firewall family ethernet-switching filter
ingress-port-voip-class-limit-tcp-icmp]
user@switch# set term icmp-connection from destination-address 192.0.2.16/28
user@switch# set term icmp-connection from protocol icmp
user@switch# set term icmp-connection then policer icmp-policer
user@switch# set term icmp-connection then count icmp-counter
user@switch# set term icmp-connection then forwarding-class best-effort
user@switch# set term icmp-connection then loss-priority high
```

7. Define the term **best-effort** with no match conditions for an implicit match on all packets that did not match any other term in the firewall filter:

```
[edit firewall family ethernet-switching filter
ingress-port-voip-class-limit-tcp-icmp]
user@switch# set term best-effort then forwarding-class best-effort
user@switch# set term best-effort then loss-priority high
```

8. Apply the firewall filter **ingress-port-voip-class-limit-tcp-icmp** as an input filter to the port interfaces for **employee-vlan** :

```
[edit interfaces]
user@switch# set ge-0/0/0 description "voice priority and tcp and icmp traffic
rate-limiting filter at ingress port"
user@switch# set ge-0/0/0 unit 0 family ethernet-switching filter input
ingress-port-voip-class-limit-tcp-icmp
user@switch# set ge-0/0/1 description "voice priority and tcp and icmp traffic
rate-limiting filter at ingress port"
user@switch# set ge-0/0/1 unit 0 family ethernet-switching filter input
ingress-port-voip-class-limit-tcp-icmp
```

9. Configure the parameters that are desired for the different schedulers.

NOTE: When you configure parameters for the schedulers, define the numbers to match your network traffic patterns.

```
[edit class-of-service]
user@switch# set schedulers voice-high buffer-size percent 15
user@switch# set schedulers voice-high priority high
user@switch# set schedulers network—control buffer-size percent 10
user@switch# set schedulers network—control priority high
user@switch# set schedulers best-effort buffer-size percent 75
user@switch# set schedulers best-effort priority low
```

10. Assign the forwarding classes to schedulers with a scheduler map:

```
[edit class-of-service]
user@switch# set scheduler-maps ethernet-diffsrv-cos-map
user@switch# set scheduler-maps ethernet-diffsrv-cos-map forwarding-class
expedited-forwarding scheduler voice-high
user@switch# set scheduler-maps ethernet-diffsrv-cos-map forwarding-class
network-control scheduler net-control
user@switch# set scheduler-maps ethernet-diffsrv-cos-map forwarding-class
best-effort scheduler best-effort
```

11. Associate the scheduler map with the outgoing interface:

```
[edit class-of-service]
user@switch# set interfaces ge—0/1/0 scheduler-map ethernet-diffsrv-cos-map
```

Results Display the results of the configuration:

```
user@switch# show
  firewall {
    policer tcp-connection-policer {
      if-exceeding {
        bandwidth-limit 1m;
        burst-size-limit 30k;
      }
      then {
        discard;
      }
    }
    policer icmp-connection-policer {
      if-exceeding {
        bandwidth-limit 1m;
        burst-size-limit 30k;
      }
      then {
        discard;
      }
    }
    family ethernet-switching {
      filter ingress-port-voip-class-limit-tcp-icmp {
        term voip-high {
          from {
            destination-mac-address 00.05.85.00.00.01;
              destination-mac-address 00.05.85.00.00.02;
            protocol udp;
          }
          then {
            forwarding-class expedited-forwarding;
            loss-priority low;
```

```
            }
        }
        term network-control {
            from {
                precedence net-control ;
            }
            then {
                forwarding-class network-control;
                loss-priority low;
            }
        }
        term tcp-connection {
            from {
                destination-address 192.0.2.16/28;
                protocol tcp;
            }
            then {
                policer tcp-connection-policer;
                    count tcp-counter;
                    forwarding-class best-effort;
                loss-priority high;
            }
        }
        term icmp-connection
            from {
                protocol icmp;
            }
            then {
                policer icmp-connection-policer;
                    count icmp-counter;
                    forwarding-class best-effort;
                loss-priority high;
            }
        }
        term best-effort {
            then {
                forwarding-class best-effort;
                loss-priority high;
            }
        }
    }
  }
}
interfaces {
  ge-0/0/0 {
    description "voice priority and tcp and icmp traffic rate-limiting filter at ingress port";
    unit 0 {
      family ethernet-switching {
        filter {
          input ingress-port-voip-class-limit-tcp-icmp;
        }
      }
    }
  }
  ge-0/0/1 {
    description "voice priority and tcp and icmp traffic rate-limiting filter at ingress port";
```

```
        unit 0 {
        family ethernet-switching {
          filter {
            input ingress-port-voip-class-limit-tcp-icmp;
          }
        }
      }
    }
  }
  scheduler-maps {
    ethernet-diffsrv-cos-map {
      forwarding-class expedited-forwarding scheduler voice-high;
        forwarding-class network-control scheduler net-control;
      forwarding-class best-effort scheduler best-effort;
    }
  }
  interfaces {
    ge/0/1/0 {
      scheduler-map ethernet-diffsrv-cos-map;
    }
  }
```

Configuring a VLAN Ingress Firewall Filter to Prevent Rogue Devices from Disrupting VoIP Traffic

To configure and apply firewall filters for port, VLAN, and router interfaces, perform these tasks:

CLI Quick Configuration

To quickly configure a VLAN firewall filter on **voice-vlan** to prevent rogue devices from using HTTP sessions to mimic the gatekeeper device that manages VoIP traffic, copy the following commands and paste them into the switch terminal window:

```
[edit]
set firewall family ethernet-switching filter ingress-vlan-rogue-block term to-gatekeeper from destination-address 192.0.2.14
set firewall family ethernet-switching filter ingress-vlan-rogue-block term to-gatekeeper from destination-port 80
set firewall family ethernet-switching filter ingress-vlan-rogue-block term to-gatekeeper then accept
set firewall family ethernet-switching filter ingress-vlan-rogue-block term from-gatekeeper from source-address 192.0.2.14
set firewall family ethernet-switching filter ingress-vlan-rogue-block term from-gatekeeper from source-port 80
set firewall family ethernet-switching filter ingress-vlan-rogue-block term from-gatekeeper then accept
set firewall family ethernet-switching filter ingress-vlan-rogue-block term not-gatekeeper from destination-port 80
set firewall family ethernet-switching filter ingress-vlan-rogue-block term not-gatekeeper then count rogue-counter
set firewall family ethernet-switching filter ingress-vlan-rogue-block term not-gatekeeper then discard
set vlans voice-vlan description "block rogue devices on voice-vlan"
set vlans voice-vlan filter input ingress-vlan-rogue-block
```

Step-by-Step Procedure

To configure and apply a VLAN firewall filter on **voice-vlan** to prevent rogue devices from using HTTP to mimic the gatekeeper device that manages VoIP traffic:

1. Define the firewall filter **ingress-vlan-rogue-block** to specify filter matching on the traffic you want to permit and restrict:

```
[edit firewall]
user@switch# set family ethernet-switching filter ingress-vlan-rogue-block
```

2. Define the term **to-gatekeeper** to accept packets that match the destination IP address of the gatekeeper:

```
[edit firewall family ethernet-switching filter ingress-vlan-rogue-block]
user@switch# set term to-gatekeeper from destination-address 192.0.2.14
user@switch# set term to-gatekeeper from destination-port 80
user@switch# set term to-gatekeeper then accept
```

3. Define the term **from-gatekeeper** to accept packets that match the source IP address of the gatekeeper:

```
[edit firewall family ethernet-switching filter ingress-vlan-rogue-block]
user@switch# set term from-gatekeeper from source-address 192.0.2.14
user@switch# set term from-gatekeeper from source-port 80
user@switch# set term from-gatekeeper then accept
```

4. Define the term **not-gatekeeper** to ensure all **voice-vlan** traffic on TCP ports is destined for the gatekeeper device:

```
[edit firewall family ethernet-switching filter ingress-vlan-rogue-block]
user@switch# set term not-gatekeeper from destination-port 80
user@switch# set term not-gatekeeper then count rogue-counter
user@switch# set term not-gatekeeper then discard
```

5. Apply the firewall filter **ingress-vlan-rogue-block** as an input filter to the VLAN interface for the VoIP telephones:

```
[edit]
user@switch# set vlans voice-vlan description "block rogue devices on voice-vlan"
user@switch# set vlans voice-vlan filter input ingress-vlan-rogue-block
```

Results

Display the results of the configuration:

```
user@switch# show
  firewall {
    family ethernet-switching {
      filter ingress-vlan-rogue-block {
        term to-gatekeeper {
          from {
            destination-address 192.0.2.14/32
            destination-port 80;
          }
          then {
            accept;
          }
        }
        term from-gatekeeper {
          from {
            source-address 192.0.2.14/32
```

```
            source-port 80;
          }
          then {
            accept;
          }
        }
        term not-gatekeeper {
          from {
            destination-port 80;
          }
          then {
            count rogue-counter;
            discard;
          }
        }
    }
  vlans {
    voice-vlan {
      description "block rogue devices on voice-vlan";
      filter {
        input ingress-vlan-rogue-block;
      }
    }
  }
```

Configuring a VLAN Firewall Filter to Count, Monitor, and Analyze Egress Traffic on the Employee VLAN

To configure and apply firewall filters for port, VLAN, and router interfaces, perform these tasks:

CLI Quick Configuration

A firewall filter is configured and applied to VLAN interfaces to filter **employee-vlan** egress traffic. Employee traffic destined for the corporate subnet is accepted but not monitored. Employee traffic destined for the Web is counted and analyzed.

To quickly configure and apply a VLAN firewall filter, copy the following commands and paste them into the switch terminal window:

```
[edit]
set firewall family ethernet-switching filter egress-vlan-watch-employee term employee-to-corp
from destination-address 192.0.2.16/28
set firewall family ethernet-switching filter egress-vlan-watch-employee term employee-to-corp
then accept
set firewall family ethernet-switching filter egress-vlan-watch-employee term employee-to-web
from destination-port 80
set firewall family ethernet-switching filter egress-vlan-watch-employee term employee-to-web
then count employee-web-counter
set firewall family ethernet-switching filter egress-vlan-watch-employee term employee-to-web
then analyzer employee-monitor
set vlans employee-vlan description "filter at egress VLAN to count and analyze employee to
Web traffic"
set vlans employee-vlan filter output egress-vlan-watch-employee
```

Step-by-Step Procedure

To configure and apply an egress port firewall filter to count and analyze **employee-vlan** traffic that is destined for the Web:

1. Define the firewall filter **egress-vlan-watch-employee**:

```
[edit firewall]
user@switch# set family ethernet-switching filter egress-vlan-watch-employee
```

2. Define the term **employee-to-corp** to accept but not monitor all **employee-vlan** traffic destined for the corporate subnet:

```
[edit firewall family ethernet-switching filter egress-vlan-watch-employee]
user@switch# set term employee-to-corp from destination-address 192.0.2.16/28
user@switch# set term employee-to-corp then accept
```

3. Define the term **employee-to-web** to count and monitor all **employee-vlan** traffic destined for the Web:

```
[edit firewall family ethernet-switching filter egress-vlan-watch-employee]
user@switch# set term employee-to-web from destination-port 80
user@switch# set term employee-to-web then count employee-web-counter
user@switch# set term employee-to-web then analyzer employee-monitor
```

NOTE: See “Example: Configuring Port Mirroring for Local Monitoring of Employee Resource Use on J-EX Series Switches” on page 3249 for information about configuring the employee-monitor analyzer.

4. Apply the firewall filter **egress-vlan-watch-employee** as an output filter to the port interfaces for the VoIP telephones:

```
[edit]
user@switch# set vlans employee-vlan description "filter at egress VLAN to count and
analyze employee to Web traffic"
user@switch# set vlans employee-vlan filter output egress-vlan-watch-employee
```

Results

Display the results of the configuration:

```
user@switch# show
  firewall {
    family ethernet-switching {
      filter egress-vlan-watch-employee {
        term employee-to-corp {
          from {
            destination-address 192.0.2.16/28
          }
          then {
            accept;
          }
        }
        term employee-to-web {
          from {
            destination-port 80;
          }
          then {
            count employee-web-counter:
```

```
            analyzer employee-monitor;
          }
        }
      }
    }
  }
  vlans {
    employee-vlan {
      description "filter at egress VLAN to count and analyze employee to Web traffic";
      filter {
        output egress-vlan-watch-employee;
      }
    }
  }
```

Configuring a VLAN Firewall Filter to Restrict Guest-to-Employee Traffic and Peer-to-Peer Applications on the Guest VLAN

To configure and apply firewall filters for port, VLAN, and router interfaces, perform these tasks:

CLI Quick Configuration

In the following example, the first filter term permits guests to talk with other guests but not employees on **employee-vlan**. The second filter term allows guests Web access but prevents them from using peer-to-peer applications on **guest-vlan**.

To quickly configure a VLAN firewall filter to restrict guest-to-employee traffic, blocking guests from talking with employees or employee hosts on **employee-vlan** or attempting to use peer-to-peer applications on **guest-vlan**, copy the following commands and paste them into the switch terminal window:

```
[edit]
set firewall family ethernet-switching filter ingress-vlan-limit-guest term guest-to-guest from
destination-address 192.0.2.33/28
set firewall family ethernet-switching filter ingress-vlan-limit-guest term guest-to-guest then
accept
set firewall family ethernet-switching filter ingress-vlan-limit-guest term
no-guest-employee-no-peer-to-peer from destination-mac-address 00.05.85.00.00.DF
set firewall family ethernet-switching filter ingress-vlan-limit-guest term
no-guest-employee-no-peer-to-peer then accept
set vlans guest-vlan description "restrict guest-to-employee traffic and peer-to-peer applications
on guest VLAN"
set vlans guest-vlan filter input ingress-vlan-limit-guest
```

Step-by-Step Procedure

To configure and apply a VLAN firewall filter to restrict guest-to-employee traffic and peer-to-peer applications on **guest-vlan**:

1. Define the firewall filter **ingress-vlan-limit-guest**:

   ```
   [edit firewall]
   set firewall family ethernet-switching filter ingress-vlan-limit-guest
   ```

2. Define the term **guest-to-guest** to permit guests on the **guest-vlan** to talk with other guests but not employees on the **employee-vlan**:

   ```
   [edit firewall family ethernet-switching filter ingress-vlan-limit-guest]
   user@switch# set term guest-to-guest from destination-address 192.0.2.33/28
   user@switch# set term guest-to-guest then accept
   ```

3. Define the term **no-guest-employee-no-peer-to-peer** to allow guests on **guest-vlan** Web access but prevent them from using peer-to-peer applications on the **guest-vlan**.

NOTE: The destination-mac-address is the default gateway, which for any host in a VLAN is the next-hop router.

```
[edit firewall family ethernet-switching filter ingress-vlan-limit-guest]
user@switch# set term no-guest-employee-no-peer-to-peer from
destination-mac-address 00.05.85.00.00.DF
user@switch# set term no-guest-employee-no-peer-to-peer then accept
```

4. Apply the firewall filter **ingress-vlan-limit-guest** as an input filter to the interface for **guest-vlan** :

```
[edit]
user@switch# set vlans guest-vlan description "restrict guest-to-employee traffic and
peer-to-peer applications on guest VLAN"
user@switch# set vlans guest-vlan filter input ingress-vlan-limit-guest
```

Results Display the results of the configuration:

```
user@switch# show
  firewall {
    family ethernet-switching {
      filter ingress-vlan-limit-guest {
        term guest-to-guest {
          from {
            destination-address 192.0.2.33/28;
          }
          then {
            accept;
          }
        }
        term no-guest-employee-no-peer-to-peer {
          from {
            destination-mac-address 00.05.85.00.00.DF;
          }
          then {
            accept;
          }
        }
      }
    }
  }
  vlans {
    guest-vlan {
      description "restrict guest-to-employee traffic and peer-to-peer applications on
        guest VLAN";
      filter {
        input ingress-vlan-limit-guest;
      }
    }
  }
```

Configuring a Router Firewall Filter to Give Priority to Egress Traffic Destined for the Corporate Subnet

To configure and apply firewall filters for port, VLAN, and router interfaces, perform these tasks:

CLI Quick Configuration

To quickly configure a firewall filter for a routed port (Layer 3 uplink module) to filter **employee-vlan** traffic, giving highest forwarding-class priority to traffic destined for the corporate subnet, copy the following commands and paste them into the switch terminal window:

```
[edit]
set firewall family inet filter egress-router-corp-class term corp-expedite from destination-address 192.0.2.16/28
set firewall family inet filter egress-router-corp-class term corp-expedite then forwarding-class expedited-forwarding
set firewall family inet filter egress-router-corp-class term corp-expedite then loss-priority low
set firewall family inet filter egress-router-corp-class term not-to-corp then accept
set interfaces ge-0/1/0 description "filter at egress router to expedite destined for corporate network"
set ge-0/1/0 unit 0 family inet address 103.104.105.1
set interfaces ge-0/1/0 unit 0 family inet filter output egress-router-corp-class
```

Step-by-Step Procedure

To configure and apply a firewall filter to a routed port (Layer 3 uplink module) to give highest priority to **employee-vlan** traffic destined for the corporate subnet:

1. Define the firewall filter **egress-router-corp-class**:

   ```
   [edit]
   user@switch# set firewall family inet filter egress-router-corp-class
   ```

2. Define the term **corp-expedite**:

   ```
   [edit firewall]
   user@switch# set family inet filter egress-router-corp-class term corp-expedite from destination-address 192.0.2.16/28
   user@switch# set family inet filter egress-router-corp-class term corp-expedite then forwarding-class expedited-forwarding
   user@switch# set family inet filter egress-router-corp-class term corp-expedite then loss-priority low
   ```

3. Define the term **not-to-corp**:

   ```
   [edit firewall]
   user@switch# set family inet filter egress-router-corp-class term not-to-corp then accept
   ```

4. Apply the firewall filter **egress-router-corp-class** as an output filter for the port on the switch's uplink module, which provides a Layer 3 connection to a router:

   ```
   [edit interfaces]
   user@switch# set ge-0/1/0 description "filter at egress router to expedite employee traffic destined for corporate network"
   user@switch# set ge-0/1/0 unit 0 family inet address 103.104.105.1
   user@switch# set ge-0/1/0 unit 0 family inet filter output egress-router-corp-class
   ```

Results

Display the results of the configuration:

```
user@switch# show
```

```
firewall {
  family inet {
    filter egress-router-corp-class {
      term corp-expedite {
        from {
          destination-address 192.0.2.16/28;
        }
        then {
          forwarding-class expedited-forwarding;
          loss-priority low;
        }
      }
      term not-to-corp {
        then {
          accept;
        }
      }
    }
  }
}
interfaces {
  ge-0/1/0 {
    unit 0 {
      description "filter at egress router interface to expedite employee traffic destined
        for corporate network";
      family inet {
        source-address 103.104.105.1
        filter {
          output egress-router-corp-class;
        }
      }
    }
  }
}
```

Verification

To confirm that the firewall filters are working properly, perform the following tasks:

- Verifying that Firewall Filters and Policers are Operational on page 2771
- Verifying that Schedulers and Scheduler-Maps are Operational on page 2772

Verifying that Firewall Filters and Policers are Operational

Purpose Verify the operational state of the firewall filters and policers that are configured on the switch.

Action Use the operational mode command:

```
user@switch> show firewall
Filter: ingress-port-voip-class-limit-tcp-icmp
Counters:
Name                                                Packets
icmp-counter                                              0
tcp-counter                                               0
Policers:
```

```
Name                                                    Packets
icmp-connection-policer                                       0
tcp-connection-policer                                        0

Filter: ingress-vlan-rogue-block

Filter: egress-vlan-watch-employee
Counters:
Name                                                    Packets
employee-web-counter                                          0
```

Meaning The **show firewall** command displays the names of the firewall filters, policers, and counters that are configured on the switch. The output fields show byte and packet counts for all configured counters and the packet count for all policers.

Verifying that Schedulers and Scheduler-Maps are Operational

Purpose Verify that schedulers and scheduler-maps are operational on the switch.

Action Use the operational mode command:

```
user@switch>  show class-of-service scheduler-map

Scheduler map: default, Index: 2

  Scheduler: default-be, Forwarding class: best-effort, Index: 20
    Transmit rate: 95 percent, Rate Limit: none, Buffer size: 95 percent,
    Priority: low
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             non-TCP         1    default-drop-profile
      Low             TCP             1    default-drop-profile
      High            non-TCP         1    default-drop-profile
      High            TCP             1    default-drop-profile

  Scheduler: default-nc, Forwarding class: network-control, Index: 22
    Transmit rate: 5 percent, Rate Limit: none, Buffer size: 5 percent,
    Priority: low
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             non-TCP         1    default-drop-profile
      Low             TCP             1    default-drop-profile
      High            non-TCP         1    default-drop-profile
      High            TCP             1    default-drop-profileScheduler map:
ethernet-diffsrv-cos-map, Index: 21657

  Scheduler: best-effort, Forwarding class: best-effort, Index: 61257
    Transmit rate: remainder, Rate Limit: none, Buffer size: 75 percent,
    Priority: low
    Drop profiles:
      Loss priority   Protocol    Index    Name
      Low             non-TCP         1    <default-drop-profile>
      Low             TCP             1    <default-drop-profile>
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

  Scheduler: voice-high, Forwarding class: expedited-forwarding, Index: 3123
    Transmit rate: remainder, Rate Limit: none, Buffer size: 15 percent,
    Priority: high
```

```
  Drop profiles:
    Loss priority   Protocol    Index    Name
    Low             non-TCP         1    <default-drop-profile>
    Low             TCP             1    <default-drop-profile>
    High            non-TCP         1    <default-drop-profile>
    High            TCP             1    <default-drop-profile>

Scheduler: net-control, Forwarding class: network-control, Index: 2451
  Transmit rate: remainder, Rate Limit: none, Buffer size: 10 percent,
  Priority: high
  Drop profiles:
    Loss priority   Protocol    Index    Name
    Low             non-TCP         1    <default-drop-profile>
    Low             TCP             1    <default-drop-profile>
    High            non-TCP         1    <default-drop-profile>
    High            TCP             1    <default-drop-profile>
```

Meaning Displays statistics about the configured schedulers and schedulers-maps.

Related Documentation

- Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on J-EX Series Switches on page 3254
- Example: Configuring CoS on J-EX Series Switches on page 2883
- Configuring Firewall Filters (CLI Procedure) on page 2779
- Configuring Firewall Filters (J-Web Procedure) on page 2784
- Configuring Policers to Control Traffic Rates (CLI Procedure) on page 2788
- Firewall Filter Match Conditions and Actions for J-EX Series Switches on page 2728
- [edit firewall] Configuration Statement Hierarchy on page 42

Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on J-EX Series Switches

Administrators can configure filter-based forwarding on a J-EX Series switch by using a firewall filter to forward matched traffic to a specific virtual routing instance.

This example describes how to set up filter-based forwarding:

Requirements

This example uses the following software and hardware components:

- One J-EX Series switch

Overview and Topology

In this example, traffic from one application server that is destined for a different application server is matched by a firewall filter based on the IP address. Any matching packets are routed to a particular virtual routing instance that first sends all traffic to a security device, then forwards it to the designated destination address.

Configuration

To configure filter-based forwarding:

CLI Quick Configuration

To quickly create and configure filter-based forwarding, copy the following commands and paste them into the switch terminal window:

```
[edit]
set interfaces ge-0/0/0 unit 0 family inet address 10.1.0.1/24
set interfaces ge-0/0/3 unit 0 family inet address 10.1.3.1/24
set firewall family inet filter fil term t1 from source-address 1.1.1.1/32
set firewall family inet filter fil term t1 from protocol tcp
set interfaces ge-0/0/0 unit 0 family inet filter input fil
set routing-instances vrf01 instance-type virtual-router
set routing-instances vrf01 interface ge-0/0/1.0
set routing-instances vrf01 interface ge-0/0/3.0
set routing-instances vrf01 routing-options static route 12.34.56.0/24 next-hop 10.1.3.254
set firewall family inet filter fil term t1 then routing-instance vrf01
```

Step-by-Step Procedure

To configure filter-based forwarding:

1. Create interfaces to the application servers:

```
[edit]
user@switch# set interfaces ge-0/0/0 unit 0 family inet address 10.1.0.1/24
user@switch# set interfaces ge-0/0/3 unit 0 family inet address 10.1.3.1/24
```

2. Create a firewall filter that matches the correct source address:

```
[edit]
user@switch# set firewall family inet filter fil term t1 from source-address 1.1.1.1/32
user@switch# set firewall family inet filter fil term t1 from protocol tcp
```

3. Associate the filter with the source application server's interface:

```
[edit]
user@switch# set interfaces ge-0/0/0 unit 0 family inet filter input fil
```

4. Create a virtual router:

```
[edit]
user@switch# set routing-instances vrf01 instance-type virtual-router
```

5. Associate the interfaces with the virtual router:

```
[edit]
user@switch# set routing-instances vrf01 interface ge-0/0/1.0
user@switch# set routing-instances vrf01 interface ge-0/0/3.0
```

6. Configure the routing information for the virtual routing instance:

```
[edit]
```

```
user@switch# set routing-instances vrf01 routing-options static route 12.34.56.0/24
next-hop 10.1.3.254
```

7. Set the filter to forward packets to the virtual router you created:

```
[edit]
user@switch# set firewall family inet filter fil term t1 then routing-instance vrf01
```

Results Check the results of the configuration:

```
user@switch> show configuration
interfaces {
    ge-0/0/0 {
        unit 0 {
            family inet {
                filter {
                    input fil;
                }
                address 10.1.0.1/24;
            }
        }
    }
    ge-0/0/3 {
        unit 0 {
            family inet {
                address 10.1.3.1/24;
            }
        }
    }
}
firewall {
    family inet {
        filter fil {
            term t1 {
                from {
                    source-address {
                        1.1.1.1/32;
                    }
                    protocol tcp;
                }
                then {
                    routing-instance vrf01;
                }
            }
        }
    }
}
    routing-instances {
        vrf01 {
            instance-type virtual-router;
            interface ge-0/0/1.0;
            interface ge-0/0/3.0;
            routing-options {
                static {
                    route 12.34.56.0/24 next-hop 10.1.3.254;
                }
```

```
    }
  }
}
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying That Filter-Based Forwarding Was Configured on page 2776

Verifying That Filter-Based Forwarding Was Configured

Purpose Verify that filter-based forwarding was properly enabled on the switch.

Action 1. Use the **show interfaces filters** command:

```
user@switch> show interfaces filters ge-0/0/0.0

Interface        Admin Link Proto Input Filter         Output Filter
ge-0/0/0.0       up    down inet  fil
```

2. Use the **show route forwarding-table** command:

```
user@switch> show route forwarding-table

Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            user     1 0:12:f2:21:cf:0    ucst   331     4 me0.0
default            perm     0                    rjct    36     3
0.0.0.0/32         perm     0                    dscd    34     1
10.1.0.0/24        ifdn     0                    rslv   613     1
ge-0/0/0.0
10.1.0.0/32        iddn     0 10.1.0.0           recv   611     1
ge-0/0/0.0
10.1.0.1/32        user     0                    rjct    36     3
10.1.0.1/32        intf     0 10.1.0.1           locl   612     2
10.1.0.1/32        iddn     0 10.1.0.1           locl   612     2
10.1.0.255/32      iddn     0 10.1.0.255         bcst   610     1
ge-0/0/0.0
10.1.1.0/26        ifdn     0                    rslv   583     1 vlan.0
10.1.1.0/32        iddn     0 10.1.1.0           recv   581     1 vlan.0
10.1.1.1/32        user     0                    rjct    36     3
10.1.1.1/32        intf     0 10.1.1.1           locl   582     2
10.1.1.1/32        iddn     0 10.1.1.1           locl   582     2
10.1.1.63/32       iddn     0 10.1.1.63          bcst   580     1 vlan.0
255.255.255.255/32 perm     0                    bcst    32     1

Routing table: vrf01.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct   559     2
0.0.0.0/32         perm     0                    dscd   545     1
10.1.3.0/24        ifdn     0                    rslv   617     1
ge-0/0/3.0
10.1.3.0/32        iddn     0 10.1.3.0           recv   615     1
ge-0/0/3.0
10.1.3.1/32        user     0                    rjct   559     2
10.1.3.1/32        intf     0 10.1.3.1           locl   616     2
10.1.3.1/32        iddn     0 10.1.3.1           locl   616     2
```

```
10.1.3.255/32      iddn     0 10.1.3.255         bcst   614     1
ge-0/0/3.0
224.0.0.0/4        perm     0                    mdsc   546     1
224.0.0.1/32       perm     0 224.0.0.1          mcst   529     1
255.255.255.255/32 perm     0                    bcst   543     1

Routing table: default.iso
ISO:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    60     1

Routing table: vrf01.iso
ISO:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct   600     1
```

Meaning The output indicates that the filter was created on the interface and that the virtual routing instance is forwarding matching traffic to the correct IP address.

Related Documentation

- Configuring Firewall Filters (CLI Procedure) on page 2779
- Configuring Static Routing (CLI Procedure) on page 1444
- Configuring Static Routing (J-Web Procedure) on page 1444
- Understanding Filter-Based Forwarding for J-EX Series Switches on page 2753

CHAPTER 102

Configuring Firewall Filters

Configuring Firewall Filters (CLI Procedure)

You configure firewall filters on J-EX Series switches to control traffic that enters ports on the switch or enters and exits VLANs on the network and Layer 3 (routed) interfaces. To configure a firewall filter you must configure the filter and then apply it to a port, VLAN, or Layer 3 interface.

Configuring a Firewall Filter

To configure a firewall filter:

1. Configure the family address type for the firewall filter:

   - For a firewall filter that is applied to a port or VLAN, specify the family address type **ethernet-switching** to filter Layer 2 (Ethernet) packets and Layer 3 (IP) packets, for example:

     ```
     [edit firewall]
     user@switch# set family ethernet-switching
     ```

   - For a firewall filter that is applied to a Layer 3 (routed) interface:

     - To filter IPv4 packets, specify the family address type **inet**, for example:

       ```
       [edit firewall]
       user@switch# set family inet
       ```

     - To filter IPv6 packets, specify the family address type **inet6**, for example:

       ```
       [edit firewall]
       user@switch# set family inet6
       ```

   

   NOTE: You can configure firewall filters for both IPv4 and IPv6 traffic on the same Layer 3 interface.

2. Specify the filter name:

   ```
   [edit firewall family ethernet-switching]
   user@switch# set filter ingress-port-filter
   ```

   The filter name can contain letters, numbers, and hyphens (-) and can have a maximum of 64 characters. Each filter name must be unique.

3. If you want to apply a firewall filter to multiple interfaces and name individual firewall counters specific to each interface, configure the **interface-specific** option:

   ```
   [edit firewall family ethernet-switching filter ingress-port-filter]
   user@switch# set interface-specific
   ```

4. Specify a term name:

   ```
   [edit firewall family ethernet-switching filter ingress-port-filter]
   user@switch# set term term-one
   ```

   The term name can contain letters, numbers, and hyphens (-) and can have a maximum of 64 characters.

   A firewall filter can contain one or more terms. Each term name must be unique within a filter.

   

   NOTE: For J-EX4200 switches, the maximum number of terms allowed per firewall filter is 2048. For J-EX8200 switches, the maximum number of terms allowed per firewall filter is 32768. If you attempt to configure a firewall filter that exceeds these limits, the switch returns an error message when you commit the configuration.

5. In each firewall filter term, specify the match conditions to use to match components of a packet.

   To specify match conditions to match on packets that contain a specific source-address and source-port—for example:

   ```
   [edit firewall family ethernet-switching filter ingress-port-filter term
   term-one]
   user@switch# set from source-address 192.0.2.14
   user@switch# set from source-port 80
   ```

   You can specify one or more match conditions in a single **from** statement. For a match to occur, the packet must match all the conditions in the term.

   The **from** statement is optional, but if included in a term, the **from** statement cannot be empty. If you omit the **from** statement, all packets are considered to match.

6. In each firewall filter term, specify the actions to take if the packet matches all the conditions in that term.

   You can specify an action and/or action modifiers:

   - To specify a filter action, for example, to discard packets that match the conditions of the filter term:

     ```
     [edit firewall family ethernet-switching filter ingress-port-filter term
     term-one]
     user@switch# set then discard
     ```

     You can specify no more than one action (**accept**, **discard**, or **routing-instance**) per filter term.

   - To specify action modifiers, for example, to count and classify packets in a forwarding class:

     ```
     [edit firewall family ethernet-switching filter ingress-port-filter term
     term-one]
     user@switch# set then count counter-one
     user@switch# set then forwarding-class expedited-forwarding
     ```

     You can specify any of the following action modifiers in a **then** statement:

     - **analyzer** ***analyzer-name***—Mirror port traffic to a specified destination port or VLAN that is connected to a protocol analyzer application. An **analyzer** must be configured under the **ethernet-switching** family address type. See “Configuring Port Mirroring to Analyze Traffic (CLI Procedure)” on page 3260.
     - **count** ***counter-name***—Count the number of packets that pass this filter term.

       NOTE: We recommend that you configure a counter for each term in a firewall filter, so that you can monitor the number of packets that match the conditions specified in each filter term.

     - **forwarding-class** ***class***—Classify packets in a forwarding class.

- **loss-priority** ***priority***—Set the priority of dropping a packet.
- **policer** ***policer-name***—Apply rate-limiting to the traffic.

If you omit the **then** statement or do not specify an action, packets that match all the conditions in the **from** statement are accepted. However, you must always explicitly configure an action and/or action modifier in the **then** statement. You can include no more than one action statement, but you can use any combination of action modifiers. For an action or action modifier to take effect, all conditions in the **from** statement must match.

NOTE: Implicit discard is also applicable to a firewall filter applied to the loopback interface, lo0.

Applying a Firewall Filter to a Port on a Switch

To apply a firewall filter to an ingress port on a switch:

1. Specify the interface name and provide a meaningful description of the firewall filter and the interface to which the filter is applied:

```
[edit interfaces]
user@switch# set ge-0/0/1 description "filter to limit tcp traffic filter at trunk port for employee-vlan and voice-vlan applied on the interface"
```

NOTE: Providing the description is optional.

2. Specify the unit number and family address type for the interface:

```
[edit interfaces]
user@switch# set ge-0/0/1 unit 0 family ethernet-switching
```

For firewall filters that are applied to ports, the family address type must be **ethernet-switching**.

3. To apply a firewall filter to filter packets that are entering a port:

```
[edit interfaces]
user@switch# set ge-0/0/1 unit 0 family ethernet-switching filter input ingress-port-filter
```

You cannot apply a firewall filter to filter packets that are exiting ports.

NOTE: You can apply no more than one firewall filter per ingress port.

Applying a Firewall Filter to a VLAN on a Network

To apply a firewall filter to a VLAN:

1. Specify the VLAN name and VLAN ID and provide a meaningful description of the firewall filter and the VLAN to which the filter is applied:

```
[edit vlans]
user@switch# set employee-vlan vlan-id 20 vlan-description "filter to rate limit traffic applied on employee-vlan"
```

NOTE: Providing the description is optional.

2. Apply firewall filters to filter packets that are entering or exiting the VLAN:

   - To apply a firewall filter to filter packets that are entering the VLAN:

```
[edit vlans]
user@switch# set employee-vlan vlan-id 20 filter input ingress-vlan-filter
```

   - To apply a firewall filter to filter packets that are exiting the VLAN:

```
[edit vlans]
user@switch# set employee-vlan vlan-id 20 filter output egress-vlan-filter
```

NOTE: You can apply no more than one firewall filter per VLAN, per direction.

Applying a Firewall Filter to a Layer 3 (Routed) Interface

To apply a firewall filter to a Layer 3 (routed) interface on a switch:

1. Specify the interface name and provide a meaningful description of the firewall filter and the interface to which the filter is applied:

```
[edit interfaces]
user@switch# set ge-0/1/0 description "filter to count and monitor employee-vlan traffic applied on layer 3 interface"
```

NOTE: Providing the description is optional.

2. Specify the unit number, family address type, and address for the interface:

```
[edit interfaces]
user@switch# set ge-0/1/0 unit 0 family inet address 10.10.10.1/24
```

   For firewall filters applied to Layer 3 (routed) interfaces, the family address type must be **inet** (for IPv4 traffic) or **inet6** (for IPv6 traffic).

3. You can apply firewall filters to filter packets that are entering or exiting a Layer 3 (routed) interface:

   - To apply a firewall filter to filter packets that are entering a Layer 3 interface:

```
[edit interfaces]
```

```
user@switch# set ge-0/1/0 unit 0 family inet address 10.10.10.1/24 filter input
ingress-router-filter
```

- To apply a firewall filter to filter packets that are exiting a Layer 3 interface:

```
[edit interfaces]
user@switch# set ge-0/1/0 unit 0 family inet address 10.10.10.1/24 filter output
egress-router-filter
```

NOTE: You can apply no more than one firewall filter per Layer 3 interface, per direction.

Related Documentation

- Configuring Firewall Filters (J-Web Procedure) on page 2784
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on J-EX Series Switches on page 2773
- Verifying That Firewall Filters Are Operational on page 2799
- Monitoring Firewall Filter Traffic on page 2800
- Configuring Policers to Control Traffic Rates (CLI Procedure) on page 2788
- Assigning Multifield Classifiers in Firewall Filters to Specify Packet-Forwarding Behavior (CLI Procedure) on page 2791
- Firewall Filter Match Conditions and Actions for J-EX Series Switches on page 2728
- Firewall Filters for J-EX Series Switches Overview on page 2721

Configuring Firewall Filters (J-Web Procedure)

You configure firewall filters on J-EX Series switches to control traffic that enters ports on the switch or enters and exits VLANs on the network and Layer 3 (routed) interfaces. To configure a firewall filter you must configure the filter and then apply it to a port, VLAN, or Layer 3 interface.

To configure firewall filter settings using the J-Web interface:

1. Select **Configure** > **Security** > **Filters**.

   The Firewall Filter Configuration page displays a list of all configured port/VLAN or router filters and the ports or VLANs associated with a particular filter.

NOTE: After you make changes to the configuration in this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See "Using the Commit Options to Commit Configuration Changes (J-Web Procedure)" on page 334 for details about all commit options.

2. Click one:
   - **Add**—Select this option to create a new filter. Enter information as specified in Table 343 on page 2785.
   - **Edit**—Select this option to edit an existing filter. Enter information as specified in Table 343 on page 2785.
   - **Delete**—Select this option to delete a filter.
   - **Term Up**—Select this option to move a term up in the filter term list.
   - **Term Down**—Select this option to move a term down in the filter term list.

Table 343: Create a New Filter

Field	Function	Your Action
Filter tab		
Filter type	Specifies the filter type: port/VLAN firewall filter or router firewall filter.	Select the filter type.
Filter name	Specifies the name for the filter.	Enter a name.
Select terms to be part of the filter	Specifies the terms to be associated with the filter. Add new terms or edit existing terms.	Click **Add** to add new terms. Enter information as specified in Table 344 on page 2786 and Table 345 on page 2786.
Association tab		
Port Associations	Specifies the ports with which the filter is associated. NOTE: For a port/VLAN filter type, only Ingress direction is supported for port association.	1. Click **Add**. 2. Select the direction: Ingress or Egress. 3. Select the ports. 4. Click **OK**.
VLAN Associations	Specifies the VLANs with which the filter is associated. NOTE: Because router firewall filters can be associated with ports only, this section is not displayed for a router firewall filter.	1. Click **Add**. 2. Select the direction: Ingress or Egress. 3. Select the VLANs. 4. Click **OK**.

Table 344: Create a New Term

Field	Function	Your Action
Term Name	Specifies the name of the term.	Enter a name.
Protocols	Specifies the protocols to be associated with the term.	1. Click **Add**. 2. Select the protocols. 3. Click **OK**.
Source	Specifies the source IP address, MAC address, and available ports. NOTE: MAC address is specified only for port/VLAN filters.	To specify the IP address, click **Add > IP** and enter the IP address. To specify the MAC address, click **Add > MAC** and enter the MAC address. To specify the ports (interfaces), click **Add > Ports** and enter the port number. To delete the IP address, MAC address, or port details, select it and click **Remove**.
Destination	Specifies the destination IP address, MAC address, and available ports. NOTE: MAC address is specified only for port/VLAN filters.	To specify the IP address, click **Add > IP** and enter the IP address. To specify the MAC address, click **Add > MAC** and enter the MAC address. To specify the ports (interfaces), click **Add > Ports** and enter the port number. To delete the IP address, MAC address, or port details, select it and click **Remove**.
Action	Specifies the packet action for the term.	Select one: • Accept • Discard
More	Specifies advanced configuration options for the filter.	Select the match conditions as specified in Table 345 on page 2786. Select the packet action for the term as specified in Table 345 on page 2786.

Table 345: Advanced Options for Terms

Table	Function	Your Action
ICMP Type	Specifies the ICMP packet type field. Typically, you specify this match condition in conjunction with the protocol match condition to determine which protocol is being used on the port.	Select the option from the list.

Table 345: Advanced Options for Terms *(continued)*

Table	Function	Your Action
ICMP Code	Specifies more specific information than ICMP type. Because the value's meaning depends upon the associated ICMP type, you must specify **icmp-type** along with **icmp-code.** The keywords are grouped by the ICMP type with which they are associated.	Select a value from the list.
DSCP	Specifies the Differentiated Services code point (DSCP). The DiffServ protocol uses the type-of-service (ToS) byte in the IP header. The most significant six bits of this byte form the DSCP.	Select the DSCP number from the list.
Precedence	Specifies IP precedence. NOTE: IP precedence and DSCP number cannot be specified together for the same term.	Select the option from the list.
IP Options	Specifies the presence of the options field in the IP header.	Select the option from the list.
Interface	Specifies the interface on which the packet is received.	Select the interface from the list.
Ether type	Specifies the Ethernet type field of a packet. NOTE: This option is not applicable for a routing filter.	Select a value from the list.
Dot 1q user priority	Specifies the user-priority field of the tagged Ethernet packet. User-priority values can be 0–7. In place of the numeric value, you can specify one of the following text synonyms (the field values are also listed) • background (1)—Background • best-effort (0)—Best effort • controlled-load (4)—Controlled load • excellent-load (3)—Excellent load • network-control (7)—Network control reserved traffic • standard (2)—Standard or Spare • video (5)—Video • voice (6)—Voice NOTE: This option is not applicable for a routing filter.	Select a value from the list.
VLAN	Specifies the VLAN to be associated with the packet. NOTE: This option is not applicable for a routing filter.	Select the VLAN from the list.
TCP Flags	Specifies one or more TCP flags. NOTE: TCP flags are supported on ingress ports, VLANs, and router interfaces.	Select the option **TCP Initial** or enter a combination of TCP flags.

Table 345: Advanced Options for Terms *(continued)*

Table	Function	Your Action
Fragmentation Flags	Specifies the IP fragmentation flags. NOTE: Fragmentation flags are supported on ingress ports, VLANs, and router interfaces.	Select either the option **is-fragment** or enter a combination of fragment action flags.
Dot1q tag	Specifies the value for tag field in the Ethernet header. Values can be from 1 through 4095. NOTE: This option is not applicable for a routing filter.	Enter the value.
Action		
Counter name	Specifies the count of the number of packets that pass this filter, term, or policer.	Enter a value.
Forwarding class	Classifies the packet into one of the following forwarding classes: • assured-forwarding • best-effort • expedited-forwarding • network-control • user-defined	Select the option from the list.
Loss priority	Specifies the packet loss priority. NOTE: Forwarding class and loss priority should be specified together for the same term.	Enter the value.
Analyzer	Specifies whether to perform port-mirroring on packets. Port-mirroring copies all packets entering one switch port to a network monitoring connection on another switch port.	Select the analyzer (port mirroring configuration) from the list.

Related Documentation

- Configuring Firewall Filters (CLI Procedure) on page 2779
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Verifying That Firewall Filters Are Operational on page 2799
- Firewall Filters for J-EX Series Switches Overview on page 2721
- Firewall Filter Match Conditions and Actions for J-EX Series Switches on page 2728

Configuring Policers to Control Traffic Rates (CLI Procedure)

You can configure policers to rate limit traffic on J-EX Series switches. After you configure a policer, you can include it in an ingress firewall filter configuration.

When you configure a firewall filter, you can specify a policer action for any term or terms within the filter. All traffic that matches a term that contains a policer action goes through

the policer that the term references. Each policer that you configure includes an implicit counter. To get term-specific packet counts, you must configure a new policer for each filter term that requires policing.

The following policer limits apply on the switch:

- A maximum of 512 policers can be configured for port firewall filters.
- A maximum of 512 policers can be configured for VLAN and Layer 3 firewall filters.

If the policer configuration exceeds these limits, the switch returns the following message after the commit operation:

```
Cannot assign policers: Max policer limit reached
```

1. Configuring Policers on page 2789
2. Specifying Policers in a Firewall Filter Configuration on page 2790
3. Applying a Firewall Filter That Is Configured with a Policer on page 2790

Configuring Policers

To configure a policer:

1. Specify the name of the policer:

   ```
   [edit firewall]
   user@switch# set policer policer-one
   ```

   The policer name can contain letters, numbers, and hyphens (-) and can be up to 64 characters long.

2. Configure rate limiting for the policer:

   a. Specify the bandwidth limit in bits per second (bps) to control the traffic rate on an interface:

      ```
      [edit firewall policer policer-one]
      user@switch# set if-exceeding bandwidth-limit 300k
      ```

      The range for the bandwidth limit is 1k through 102.3g bps.

   b. Specify the maximum allowed burst size to control the amount of traffic bursting:

      ```
      [edit firewall policer policer-one]
      user@switch# set if-exceeding burst-size-limit 500k
      ```

      To determine the value for the burst-size limit, multiply the bandwidth of the interface on which the filter is applied by the amount of time to allow a burst of traffic at that bandwidth to occur:

      burst size = bandwidth * allowable time for burst traffic

      The range for the burst-size limit is 1 through 2,147,450,880 bytes.

3. Specify the policer action **discard** to discard packets that exceed the rate limits:

   ```
   [edit firewall policer]
   user@switch# set policer-one then discard
   ```

Discard is the only supported policer action.

Specifying Policers in a Firewall Filter Configuration

To reference a policer for a single firewall, configure a filter term that includes the policer action:

```
[edit firewall family ethernet-switching]
user@switch# set filter limit-hosts term term-one from source-address 192.0.2.16/28
user@switch# set filter limit-hosts term term-one then policer policer-one
```

Applying a Firewall Filter That Is Configured with a Policer

A firewall filter that is configured with one or more policer actions, like any other filter, must be applied to a port, VLAN, or Layer 3 interface. For information about applying firewall filters, see the sections on applying firewall filters in “Configuring Firewall Filters (CLI Procedure)” on page 2779.

NOTE: You can include policer actions on ingress firewall filters only.

Related Documentation

- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Configuring Firewall Filters (CLI Procedure) on page 2779
- Configuring Firewall Filters (J-Web Procedure) on page 2784
- Verifying That Policers Are Operational on page 2800
- Understanding the Use of Policers in Firewall Filters on page 2752

Assigning Multifield Classifiers in Firewall Filters to Specify Packet-Forwarding Behavior (CLI Procedure)

You can configure firewall filters with multifield classifiers to classify packets transiting a port, VLAN, or Layer 3 interface on a J-EX Series switch.

You specify multifield classifiers in a firewall filter configuration to set the forwarding class and packet loss priority (PLP) for incoming or outgoing packets. By default, the data traffic that is not classified is assigned to the **best-effort** class associated with queue 0.

You can specify any of the following default forwarding classes:

Forwarding class	Queue
best-effort	0
assured-forwarding	1
expedited-forwarding	5
network-control	7

To assign multifield classifiers in firewall filters:

1. Configure the family name and filter name for the filter at the **[edit firewall]** hierarchy level, for example:

```
[edit firewall]
user@switch# set family ethernet-switching
user@switch# set family ethernet-switching filter ingress-filter
```

2. Configure the terms of the filter, including the **forwarding-class** and **loss-priority** action modifiers as appropriate. When you specify a forwarding class you must also specify the packet loss priority. For example, each of the following terms examines different packet header fields and assigns an appropriate classifier and the packet loss priority:

   - The term **voice-traffic** matches packets on the **voice-vlan** and assigns the forwarding class **expedited-forwarding** and packet loss priority **low**:

```
[edit firewall family ethernet-switching filter ingress-filter]
user@switch# set term voice-traffic from vlan-id voice-vlan
user@switch# set term voice-traffic then forwarding-class expedited-forwarding
user@switch# set term voice-traffic then loss-priority low
```

   - The term **data-traffic** matches packets on **employee-vlan** and assigns the forwarding class **assured-forwarding** and packet loss priority **low**:

```
[edit firewall family ethernet-switching filter ingress-filter]
user@switch# set term data-traffic from vlan-id employee-vlan
user@switch# set term data-traffic then forwarding-class assured-forwarding
user@switch# set term data-traffic then loss-priority low
```

- Because loss of network-generated packets can jeopardize proper network operation, delay is preferable to discard of packets. The following term, **network-traffic**, assigns the forwarding class **network-control** and packet loss priority **low**:

```
[edit firewall family ethernet-switching filter ingress-filter]
user@switch# set term network-traffic from precedence net-control
user@switch# set term network-traffic then forwarding-class network
user@switch# set term network-traffic then loss-priority low
```

- The last term **accept-traffic** matches any packets that did not match on any of the preceding terms and assigns the forwarding class **best-effort** and packet loss priority **low**:

```
[edit firewall family ethernet-switching filter ingress-filter]
user@switch# set term accept-traffic from precedence net-control
user@switch# set term accept-traffic then forwarding-class best-effort
user@switch# set term accept-traffic then loss-priority low
```

3. Apply the filter **ingress-filter** to a port, VLAN or Layer 3 interface. For information about applying the filter, see "Configuring Firewall Filters (CLI Procedure)" on page 2779.

Related Documentation

- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Verifying That Firewall Filters Are Operational on page 2799
- Monitoring Firewall Filter Traffic on page 2800
- Defining CoS Classifiers (CLI Procedure) on page 2914
- Defining CoS Classifiers (J-Web Procedure) on page 2916
- Configuring Firewall Filters (CLI Procedure) on page 2779
- Configuring Firewall Filters (J-Web Procedure) on page 2784

Configuring Routing Policies (J-Web Procedure)

All routing protocols use the Junos OS routing table to store the routes that they learn and to determine which routes are advertised in the protocol packets. Routing policy allows you to control which routes the routing protocols store in and retrieve from the routing table on the routing device.

To configure routing policies for a J-EX Series switch using the J-Web interface:

1. Select **Configure** > **Routing** > **Policies**.

NOTE: After you make changes to the configuration in this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select Commit Options > Commit. See "Using the Commit Options to Commit Configuration Changes (J-Web Procedure)" on page 334 for details about all commit options.

2. Click one:
   - **Global Options**—Configures global options for policies. Enter information into the configuration page as described in Table 346 on page 2793.
   - **Add**—Configures a new policy. Select **New** and specify a policy name. To add terms, enter information into the configuration page as described in Table 347 on page 2794 . Select **Clone** to create a copy of an existing policy.
   - **Edit**—Edits an existing policy. To modify an existing term, enter information into the configuration page as described in Table 347 on page 2794.
   - **Term Up**—Moves a term up in the list.
   - **Term Down**—Moves a term down in the list.
   - **Delete**—Deletes the selected policy.
   - **Test Policy**—Tests the policy. Use this option to check whether the policy produces the results that you expect.

Table 346: Policies Global Configuration Parameters

Field	Function	Your Action
Prefix List	Specifies a list of IPv4 address prefixes for use in a routing policy statement.	To add a prefix list: 1. Click **Add**. 2. Enter a name for the prefix list. 3. To add an IP address, click **Add**. 4. Enter the IP address and the subnet mask and click **OK**. 5. Click **OK**. To edit a prefix list, click **Edit**. Edit the settings and click **OK**. To delete a prefix list, select it and click **Delete**.
BGP Community	Specifies a BGP community.	To add a BGP community: 1. Click **Add**. 2. Enter a name for the community. 3. To add a community, click **Add**. 4. Enter the community ID and click **OK**. 5. Click **OK**. To edit a BGP community, click **Edit**. Edit the settings and click **OK**. To delete a BGP community, select it and click **Delete**.

Table 346: Policies Global Configuration Parameters *(continued)*

Field	Function	Your Action
AS Path	Specifies an AS path. This is applicable to BGP only.	To add an AS path: 1. Click **Add**. 2. Enter the AS path name. 3. Enter the regular expression and click **OK**. 4. Click **OK**. To edit an AS path, click **Edit**. Edit the settings and click **OK**. To delete an AS path, select it and click **Delete**.

Table 347: Terms Configuration Parameters

Field	Function	Your Action
Term Name	Specifies a term name.	Type or select and edit the name.
Source tab		
Family	Specifies an address family protocol.	Select a value from the list.
Routing Instance	Specifies a routing instance.	Select a value from the list.
RIB	Specifies the name of a routing table.	Select a value from the list
Preference	Specifies the individual preference value for the route.	Type or select and edit the value.
Metric	Specifies a metric value. You can specify up to four metric values.	Type or select and edit the value.
Interface	Specifies a name or IP address of one or more routing device interfaces. Do not use this qualifier with protocols that are not interface-specific, such as internal BGP (IBGP).	To add an interface, select **Add > Interface**. Select the interface from the list. To add an address, select **Add > Address**. Select the address from the list. To remove an interface, select it and click **Remove**.
Prefix List	Specifies a named list of IP addresses. You can specify an exact match with incoming routes.	Click **Add**. Select the prefix list from the list and click **OK**. To remove a prefix list, select it and click **Remove**.
Protocol	Specifies the name of the protocol from which the route was learned or to which the route is being advertised.	Click **Add** and select the protocol from the list. To remove a protocol, select it and click **Remove**.

Table 347: Terms Configuration Parameters *(continued)*

Field	Function	Your Action
Policy	Specifies the name of a policy to evaluate as a subroutine.	Click **Add**. Select the policy from the list. To remove a policy, select it and click **Remove**.
More	Specifies advanced configuration options for policies.	Click **More** for advanced configuration.
OSPF Area ID	Specifies the area identifier.	Type the IP address.
BGP Origin	Specifies the origin of the AS path information.	Select a value from the list.
Local Preference	Specifies the BGP local preference.	Type a value.
Route	Specifies the type of route.	Select **External**. Select the OSPF type from the list.
AS Path	Specifies the name of an AS path regular expression.	Click **Add**. Select the AS path from the list.
Community	Specifies the name of one or more communities.	Click **Add**. Select the community from the list.
Destination tab		
Family	Specifies an address family protocol.	Select a value from the list.
Routing Instance	Specifies a routing instance.	Select a value from the list.
RIB	Specifies the name of a routing table.	Select a value from the list.
Preference	Specifies the individual preference value for the route.	Type a value.
Metric	Specifies a metric value.	Type a value.
Interface	Specifies a name or IP address of one or more routing device interfaces. Do not use this qualifier with protocols that are not interface-specific, such as internal BGP (IBGP).	To add an interface, select **Add > Interface**. Select the interface from the list. To add an address, select **Add > Address**. Select the address from the list. To delete an interface, select it and click **Remove**.
Protocol	Specifies the name of the protocol from which the route was learned or to which the route is being advertised.	Click **Add** and select the protocol from the list. To delete a protocol, select it and click **Remove**.
Action tab		

Table 347: Terms Configuration Parameters *(continued)*

Field	Function	Your Action
Action	Specifies the action to take if the conditions match.	Select a value from the list.
Default Action	Specifies that any action that is intrinsic to the protocol is overridden. This action is also nonterminating, so that various policy terms can be evaluated before the policy is terminated.	Select a value from the list.
Next	Specifies the default control action if a match occurs, and there are no further terms in the current routing policy.	Select a value from the list.
Priority	Specifies a priority for prefixes included in an OSPF import policy. Prefixes learned through OSPF are installed in the routing table based on the priority assigned to the prefixes.	Select a value from the list.
BGP Origin	Specifies the BGP origin attribute.	Select a value from the list.
AS Path Prepend	Affixes an AS number at the beginning of the AS path. The AS numbers are added after the local AS number has been added to the path. This action adds an AS number to AS sequences only, not to AS sets. If the existing AS path begins with a confederation sequence or set, the affixed AS number is placed within a confederation sequence. Otherwise, the affixed AS number is placed with a nonconfederation sequence.	Enter a value.
AS Path Expand	Extracts the last AS number in the existing AS path and affixes that AS number to the beginning of the AS path *n* times, where *n* is a number from 1 through 32. The AS number is added before the local AS number has been added to the path. This action adds AS numbers to AS sequences only, not to AS sets. If the existing AS path begins with a confederation sequence or set, the affixed AS numbers are placed within a confederation sequence. Otherwise, the affixed AS numbers are placed within a nonconfederation sequence. This option is typically used in non-IBGP export policies.	Select the type and type a value.
Load Balance Per Packet	Specifies that all next-hop addresses in the forwarding table must be installed and have the forwarding table perform per-packet load balancing. This policy action allows you to optimize VPLS traffic flows across multiple paths.	Select the check box to enable the option.
Tag	Specifies the tag value. The tag action sets the 32-bit tag field in OSPF external link-state advertisement (LSA) packets.	Select the action and type a value.
Metric	Changes the metric (MED) value by the specified negative or positive offset. This action is useful only in an external BGP (EBGP) export policy.	Select the action and type a value.
Route	Specifies whether the route is external.	Select the **External** check box to enable the option, and select the OSPF type.
Preference	Specifies the preference value.	Select the preference action and type a value.

Table 347: Terms Configuration Parameters *(continued)*

Field	Function	Your Action
Local Preference	Specifies the BGP local preference attribute.	Select the action and type a value.
Class of Service	Specifies and applies the class-of-service parameters to routes installed into the routing table. • Source class The value entered here maintains the packet counts for a route passing through your network, based on the source address. • Destination class The value entered here maintains packet counts for a route passing through your network, based on the destination address in the packet. • Forwarding class	Type the source class. Type the destination class. Type the forwarding class.

Related Documentation

- Configuring BGP Sessions (J-Web Procedure) on page 1431
- Configuring an OSPF Network (J-Web Procedure) on page 1435
- Configuring a RIP Network (J-Web Procedure) on page 1439
- Configuring Static Routing (J-Web Procedure) on page 1444
- Layer 3 Protocols Supported on J-EX Series Switches on page 13

CHAPTER 103

Verifying Firewall Filter Configuration

Verifying That Firewall Filters Are Operational

Purpose After you configure and apply firewall filters to ports, VLANs, or Layer 3 interfaces, you can perform the following task to verify that the firewall filters configured on J-EX Series switches are working properly.

Action Use the operational mode command to verify that the firewall filters on the switch are working properly:

```
user@switch> show firewall
Filter: egress-vlan-watch-employee
Counters:
Name                                                Bytes              Packets
counter-employee-web                                    0                    0
Filter: ingress-port-voip-class-limit-tcp-icmp
Counters:
Name                                                Bytes              Packets
icmp-counter                                            0                    0
Policers:
Name                                               Packets
icmp-connection-policer                                 0
tcp-connection-policer                                  0
Filter: ingress-vlan-rogue-block
Filter: ingress-vlan-limit-guest
```

Meaning The **show firewall** command displays the names of all firewall filters, policers, and counters that are configured on the switch. For each counter that is specified in a filter configuration, the output field shows the byte count and packet count for the term in which the counter is specified. For each policer that is specified in a filter configuration, the output field shows the packet count for packets that exceed the specified rate limits.

Related Documentation

- Configuring Firewall Filters (CLI Procedure) on page 2779
- Configuring Firewall Filters (J-Web Procedure) on page 2784
- Configuring Policers to Control Traffic Rates (CLI Procedure) on page 2788

- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Monitoring Firewall Filter Traffic on page 2800

Verifying That Policers Are Operational

Purpose After you configure policers and include them in firewall filter configurations, you can perform the following tasks to verify that the policers configured on J-EX Series switches are working properly.

Action Use the operational mode command to verify that the policers on the switch are working properly:

```
user@switch>  show policer
Filter: egress-vlan-watch-employee
Filter: ingress-port-filter
Filter: ingress-port-voip-class-limit-tcp-icmp
Policers:
Name                                                Packets
icmp-connection-policer                                   0
tcp-connection-policer                                    0
Filter: ingress-vlan-rogue-block
Filter: ingress-vlan-limit-guest
```

Meaning The **show policer** command displays the names of all firewall filters and policers that are configured on the switch. For each policer that is specified in a filter configuration, the output field shows the current packet count for all packets that exceed the specified rate limits.

Related Documentation

- Configuring Policers to Control Traffic Rates (CLI Procedure) on page 2788
- Configuring Firewall Filters (CLI Procedure) on page 2779
- Configuring Firewall Filters (J-Web Procedure) on page 2784
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Monitoring Firewall Filter Traffic on page 2800

Monitoring Firewall Filter Traffic

You can monitor firewall filter traffic on J-EX Series switches.

- Monitoring Traffic for All Firewall Filters and Policers That Are Configured on the Switch on page 2801
- Monitoring Traffic for a Specific Firewall Filter on page 2801
- Monitoring Traffic for a Specific Policer on page 2801

Monitoring Traffic for All Firewall Filters and Policers That Are Configured on the Switch

Purpose Perform the following task to monitor the number of packets and bytes that matched the firewall filters and monitor the number of packets that exceeded policer rate limits:

Action Use the operational mode command:

```
user@switch> show firewall
Filter: egress-vlan-watch-employee
Counters:
Name                                                Bytes              Packets
counter-employee-web                                 3348                   27
Filter: ingress-port-voip-class-limit-tcp-icmp
Counters:
Name                                                Bytes              Packets
icmp-counter                                         4100                   49
Policers:
Name                                               Packets
icmp-connection-policer                                  0
tcp-connection-policer                                   0
Filter: ingress-vlan-rogue-block
Filter: ingress-vlan-limit-guest
```

Meaning The **show firewall** command displays the names of all firewall filters, policers, and counters that are configured on the switch. The output fields show byte and packet counts for counters and packet count for policers.

Monitoring Traffic for a Specific Firewall Filter

Purpose Perform the following task to monitor the number of packets and bytes that matched a firewall filter and monitor the number of packets that exceeded the policer rate limits.

Action Use the operational mode command:

```
user@switch> show firewall filter ingress-vlan-rogue-block
Filter: ingress-vlan-rogue-block
Counters:
Name                                                Bytes              Packets
rogue-counter                                        2308                   20
```

Meaning The **show firewall filter *filter-name*** command displays the name of the firewall filter, the packet and byte count for all counters configured with the filter, and the packet count for all policers configured with the filter.

Monitoring Traffic for a Specific Policer

Purpose Perform the following task to monitor the number of packets that exceeded policer rate limits:

Action Use the operational mode command:

```
user@switch> show policer tcp-connection-policer
Filter: ingress-port-voip-class-limit-tcp-icmp
Policers:
```

```
Name                                                  Packets
tcp-connection-policer                                      0
```

Meaning The **show policer** ***policer-name*** command displays the name of the firewall filter that specifies the policer-action and displays the number of packets that exceeded rate limits for the specified filter.

Related Documentation

- Configuring Firewall Filters (CLI Procedure) on page 2779
- Configuring Firewall Filters (J-Web Procedure) on page 2784
- Configuring Policers to Control Traffic Rates (CLI Procedure) on page 2788
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Verifying That Firewall Filters Are Operational on page 2799

CHAPTER 104

Troubleshooting Firewall Filters

Troubleshooting Firewall Filters

Firewall Filter Configuration Returns a No Space Available in TCAM Message

Problem When a firewall filter configuration exceeds the amount of available ternary content addressable memory (TCAM) space, the switch returns the following **syslogd** message:

```
No space available in tcam.
Rules for filter filter-name will not be installed.
```

The switch returns this message during the commit operation if the firewall filter that has been applied to a port, VLAN, or Layer 3 interface exceeds the amount of available TCAM space. However, the commit operation for the firewall filter configuration is completed in the CLI module.

Solution When a firewall filter configuration exceeds the amount of available TCAM table space, you must configure a new firewall filter with fewer filter terms so that the space requirements for the filter do not exceed the available space in the TCAM table.

You can perform either of the following procedures to correct the problem:

To delete the firewall filter and its bind points and apply the new smaller firewall filter to the same bind points:

1. Delete the firewall filter configuration and the bind points to ports, VLANs, or Layer 3 interfaces—for example:

```
[edit]
user@switch# delete firewall family ethernet-switching filter filter-ingress-vlan
user@switch# delete vlans voice-vlan description "filter to block rogue devices on
voice-vlan"
user@switch# delete vlans voice-vlan filter input mini-filter—ingress-vlan
```

2. Commit the operation:

```
[edit]
user@switch# commit
```

3. Configure a smaller filter with fewer terms that does not exceed the amount of available TCAM space on the switch—for example:

```
[edit]
user@switch# set firewall family ethernet-switching filter new—filter-ingress-vlan ...
```

4. Apply (bind) the new firewall filter to a port, VLAN , or Layer 3 interface—for example:

```
[edit]
user@switch# set vlans voice-vlan description "filter to block rogue devices on
voice-vlan"
user@switch# set vlans voice-vlan filter input new-filter—ingress-vlan
```

5. Commit the operation:

```
[edit]
user@switch# commit
```

To apply a new firewall filter and overwrite the existing bind points:

1. Configure a firewall filter with fewer terms than the original filter:

```
[edit]
user@switch# set firewall family ethernet-switching filter new-filter-ingress-vlan...
```

2. Apply the firewall filter to the port, VLAN, or Layer 3 interfaces to overwrite the bind points of the original filter—for example:

```
[edit]
user@switch# set vlans voice-vlan description "smaller filter to block rogue devices on
voice-vlan"
user@switch# set vlans voice-vlan filter input new-filter-ingress-vlan
```

3. Commit the operation:

```
[edit]
user@switch# commit
```

Only the original bind points, and not the original firewall filter itself, are deleted.

Related Documentation

- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Verifying That Firewall Filters Are Operational on page 2799
- Configuring Firewall Filters (CLI Procedure) on page 2779
- Configuring Firewall Filters (J-Web Procedure) on page 2784

CHAPTER 105

Configuration Statements for Firewall Filters

[edit firewall] Configuration Statement Hierarchy

```
firewall {
  family family-name {
    filter filter-name {
      interface-specific;
      term term-name {
        from {
          match-conditions;
        }
        then {
          action;
          action-modifiers;
        }
      }
    }
  }
  policer policer-name {
    filter-specific;
    if-exceeding {
      bandwidth-limit bps;
      burst-size-limit bytes;
    }
    then {
      policer-action;
    }
  }
}
```

Related Documentation

- Firewall Filter Configuration Statements Supported by Junos OS for J-EX Series Switches on page 2806
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755

- Configuring Firewall Filters (CLI Procedure) on page 2779
- Configuring Policers to Control Traffic Rates (CLI Procedure) on page 2788
- Firewall Filters for J-EX Series Switches Overview on page 2721

Firewall Filter Configuration Statements Supported by the Junos OS for J-EX Series Switches

You configure firewall filters to filter packets based on their components and to perform an action on packets that match the filter.

Table 348 on page 2806 lists the options that are supported for the firewall statement in Junos OS for J-EX Series switches.

Table 348: Supported Options for Firewall Filter Statements

Statement and Option	Description
family *family-name* { }	The ***family-name*** option specifies the version or type of addressing protocol: • **any**—Filter packets based on protocol-independent match conditions. • **ethernet-switching**—Filter Layer 2 (Ethernet) packets and Layer 3 (IP) packets • **inet**—Filter IPv4 packets • **inet6**—Filter IPv6 packets
filter *filter-name* { }	The ***filter-name*** option identifies the filter. The name can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose the name in quotation marks (" ").
interface-specific	The **interface-specific** statement configures unique names for individual firewall counters specific to each interface.
term *term-name* { }	The ***term-name*** option identifies the term. The name can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose the entire name in quotation marks (" "). Each term name must be unique within a filter.

from {
match-conditions;
}
	The **from** statement is optional. If you omit it, all packets are considered to match.

then {
action;
action-modifiers;
}
	For information about the ***action*** and ***action-modifiers*** options, see "Firewall Filter Match Conditions and Actions for J-EX Series Switches" on page 2728.

Table 348: Supported Options for Firewall Filter Statements *(continued)*

Statement and Option	Description
policer *policer-name* { }	The ***policer-name*** option identifies the policer. The name can contain letters, numbers, and hyphens (-) and can be up to 64 characters long. To include spaces in the name, enclose the name in quotation marks (" ").
filter-specific	The **filter-specific** statement configures policers and counters for a specific filter name.

if-exceeding {
bandwidth-limit *bps*
burst-size-limit *bytes*
}
	The **bandwidth-limit *bps*** option specifies the traffic rate in bits per second (bps).
	
	You can specify ***bps*** as a decimal value or as a decimal number followed by one of the following abbreviations:
	
	• k (thousand)
	• m (million)
	• g (billion, which is also called a thousand million)
	
	Range: 1000 (1k) through 102,300,000,000 (102.3g) bps
	
	The **burst-size-limit *bytes*** option specifies the maximum allowed burst size to control the amount of traffic bursting. To determine the value for the burst-size limit, you can multiply the bandwidth of the interface on which the filter is applied by the amount of time (in seconds) to allow a burst of traffic at that bandwidth to occur:
	
	burst size = bandwidth * allowable time for burst traffic
	
	You can specify a decimal value or a decimal number followed by k (thousand) or m (million).
	
	Range: 1 through 2,147,450,880 bytes

then {
policer-action
}
	Use the ***policer-action*** option to specify **discard** to discard traffic that exceeds the rate limits.

The Junos OS for J-EX Series switches does not support some of the firewall filter statements that are supported by other Junos OS packages. Table 349 on page 2808 shows the firewall filter statements that are not supported by Junos OS for J-EX Series switches.

Table 349: Firewall Filter Statements That Are Not Supported by the Junos OS for J-EX Series Switches

Statements Not Supported	Statement Hierarchy Level

• interface-set *interface-set-name* {
}
• load-balance-group *group-name* {
}
• three-color-policer *name* {
}
• logical-interface-policer;
• single-rate {
}
• two-rate {
}
	[edit firewall]

• prefix-action *name* {
}
• prefix-policer {
}
• service-filter *filter-name* {
}
• simple-filter *simple-filter-name* {
}
	[edit firewall family *family-name*]

• accounting-profile *name*;	[edit firewall family *family-name* filter *filter-name*]
• logical-bandwidth-policer; • logical-interface-policer;	[edit firewall policer *policer-name*]
bandwidth-percent *number*;	[edit firewall policer *policer-name* if-exceeding]

Related Documentation

- Firewall Filter Match Conditions and Actions for J-EX Series Switches on page 2728
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Configuring Firewall Filters (CLI Procedure) on page 2779
- Configuring Policers to Control Traffic Rates (CLI Procedure) on page 2788
- Firewall Filters for J-EX Series Switches Overview on page 2721

apply-path

Syntax apply-path *path*;

Hierarchy Level [edit logical-systems *logical-system-name* policy-options **prefix-list** *name*],
[edit policy-options **prefix-list** *name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Expand a prefix list to include all prefixes pointed to by a defined path.

Options ***path***—String of elements composed of identifiers or configuration keywords that points to a set of prefixes. You can include wildcards (enclosed in angle brackets) to match more than one identifier. You cannot add a path element, including wildcards, after a leaf statement. Path elements, including wildcards, can only be used after a container statement.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Prefix Lists

as-path

Syntax as-path *name regular-expression*;

Hierarchy Level [edit dynamic policy-options],
[edit logical-systems *logical-system-name* policy-options],
[edit policy-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define an autonomous system (AS) path regular expression for use in a routing policy match condition.

Options ***name***—Name that identifies the regular expression. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (“ ”).

regular-expression—One or more regular expressions used to match the AS path.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring AS Path Regular Expressions to Use as Routing Policy Match Conditions
- Configuring Routing Policies and Policy Objects in the Dynamic Database
- **dynamic-db on page 2817**

as-path-group

Syntax

```
as-path-group group-name {
    as-path name regular-expression;
}
```

Hierarchy Level

[edit dynamic policy-options],
[edit logical-systems *logical-system-name* policy-options],
[edit policy-options]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Define a group containing multiple AS path regular expressions for use in a routing policy match condition.

Options

group-name—Name that identifies the AS path group. One or more AS path regular expressions must be listed below the **as-path-group** hierarchy.

name—Name that identifies the regular expression. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (" ").

regular-expression—One or more regular expressions used to match the AS path.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring AS Path Regular Expressions to Use as Routing Policy Match Conditions
- Configuring Routing Policies and Policy Objects in the Dynamic Database
- **dynamic-db on page 2817**

bandwidth-limit

Syntax bandwidth-limit *bps*;

Hierarchy Level [edit **firewall policer** *policer-name* **if-exceeding**]
[edit logical-systems *logical-system-name* firewall policer *policer-name* if-exceeding]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the traffic rate in bits per second.

Options ***bps*** —Traffic rate to be specified in bits per second. Specify ***bps*** as a decimal value or as a decimal number followed by one of the following abbreviations:

- k (thousand)
- m (million)
- g (billion, which is also called a thousand million)

Range:

- 1000 (1k) through 102,300,000,000 (102.3g) bps (J-EX Series switches)
- 8000 (8k) through 40,000,000,000 (40g) bps (routers)

Required Privilege Level firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Configuring Policers to Control Traffic Rates (CLI Procedure) on page 2788
- Understanding the Use of Policers in Firewall Filters on page 2752
- Rate Limiting
- Single-Rate Two-Color Policer Overview
- Configuring a Single-Rate Two-Color Policer

burst-size-limit

Syntax burst-size-limit *bytes*;

Hierarchy Level [edit **firewall policer** *policer-name* **if-exceeding**]
[edit logical-systems *logical-system-name* firewall policer *policer-name* if-exceeding]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the maximum allowed burst size to control the amount of traffic bursting.

Options ***bytes*** —Decimal value or a decimal number followed by k (thousand) or m (million).

Range:

- 1 through 2,147,450,880 bytes (J-EX Series switches)
- 1500 through 1,00,000,000,000 bytes (routers)

Required Privilege Level firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Configuring Policers to Control Traffic Rates (CLI Procedure) on page 2788
- Understanding the Use of Policers in Firewall Filters on page 2752
- Rate Limiting
- Single-Rate Two-Color Policer Overview
- Configuring a Single-Rate Two-Color Policer

community

Syntax

```
community name {
  invert-match;
  members [ community-ids ];
}
```

Hierarchy Level

[edit dynamic policy-options],
[edit logical-systems *logical-system-name* policy-options],
[edit policy-options]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Define a community or extended community for use in a routing policy match condition.

Options

name—Name that identifies the regular expression. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters. To include spaces in the name, enclose it in quotation marks (“ ”).

invert-match—Invert the results of the community expression matching.

members *community-ids*—One or more community members. If you specify more than one member, you must enclose all members in brackets.

The format for ***community-ids*** is:

as-number:*community-value*

as-number is the AS number and can be a value in the range from 0 through 65,535. ***community-value*** is the community identifier and can be a number in the range from 0 through 65,535.

You also can specify ***community-ids*** for communities as one of the following well-known community names, which are defined in RFC 1997, *BGP Communities Attribute*:

- **no-export**—Routes containing this community name are not advertised outside a BGP confederation boundary.
- **no-advertise**—Routes containing this community name are not advertised to other BGP peers.
- **no-export-subconfed**—Routes containing this community name are not advertised to external BGP peers, including peers in other members' ASs inside a BGP confederation.

You can explicitly exclude BGP community information with a static route using the **none** option. Include **none** when configuring an individual route in the **route** portion of the **static** statement to override a **community** option specified in the **defaults** portion of the statement.

The format for extended ***community-ids*** is the following:

type:*administrator*:*assigned-number*

type is the type of extended community and can be either a **bandwidth**, **target**, **origin**, **domain-id**, **src-as**, or **rt-import** community or a 16-bit number that identifies a specific BGP extended community. The **target** community identifies the destination to which the route is going. The **origin** community identifies where the route originated. The **domain-id** community identifies the OSPF domain from which the route originated. The **src-as** community identifies the autonomous system from which the route originated. The **rt-import** community identifies the route to install in the routing table.

NOTE: For **src-as**, you can specify only an AS number and not an IP address. For **rt-import**, you can specify only an IP address and not an AS number.

administrator is the administrator. It is either an AS number or an IPv4 address prefix, depending on the type of extended community.

assigned-number identifies the local provider.

The format for linking a bandwidth with an AS number is:

bandwidth:*as-number*:*bandwidth*

as-number specifies the AS number and ***bandwidth*** specifies the bandwidth in bytes per second.

NOTE: You can specify 4-byte AS numbers as defined in RFC 4893, *BGP Support for Four-octet AS Number Space*, as well as the 2-byte AS numbers. In plain-number format, you can configure a value in the range from 1 through 4,294,967,295. To configure a **target** or **origin** extended community that includes a 4-byte AS number in the plain-number format, append the letter "L" to the end of number. For example, a target community with the 4-byte AS number 334,324 and an assigned number of 132 is represented as **target:334324L:132**.

You can also use AS-dot notation when defining a 4-byte AS number for the **target** and **origin** extended communities. Specify two integers joined by a period: *16-bit high-order value in decimal*.*16-bit low-order value in decimal*. For example, the 4-byte AS number represented in plain-number format as 65546 is represented in AS-dot notation as 1.10.

For more information about configuring AS numbers, see the *Junos OS Routing Protocols Configuration Guide*.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Overview of BGP Communities and Extended Communities as Routing Policy Match Conditions
- Defining BGP Communities and Extended Communities for Use in Routing Policy Match Conditions
- Configuring Routing Policies and Policy Objects in the Dynamic Database
- **dynamic-db on page 2817**

condition

Syntax

```
condition condition-name {
    if-route-exists address table table-name;
}
```

Hierarchy Level

[edit dynamic policy-options],
[edit logical-systems *logical-system-name* policy-options],
[edit policy-options]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Define a policy condition based on the existence of routes in specific tables for use in BGP export policies.

Options

if-route-exists *address*—Specify the address of the route in question.

table *table-name*—Specify a routing table.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Routing Policy Match Conditions Based on Routing Table Entries
- Configuring Routing Policies and Policy Objects in the Dynamic Database
- **dynamic-db on page 2817**

damping

Syntax

```
damping name {
  disable;
  half-life minutes;
  max-suppress minutes;
  reuse number;
  suppress number;
}
```

Hierarchy Level [edit logical-systems *logical-system-name* policy-options],
[edit policy-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.
.

Description Define route flap damping properties to set on BGP routes.

Options

disable—Disable damping on a per-prefix basis. Any damping state that is present in the routing table for a prefix is deleted if damping is disabled.

half-life *minutes*—Decay half-life. ***minutes*** is the interval after which the accumulated figure-of-merit value is reduced by half if the route remains stable.

Range: 1 through 45

Default: 15 minutes

max-suppress *minutes*—Maximum hold-down time. ***minutes*** is the maximum time that a route can be suppressed no matter how unstable it has been.

Range: 1 through 720

Default: 60 minutes

name—Name that identifies the set of damping parameters. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (“ ”).

reuse *number*—Reuse threshold. ***number*** is the figure-of-merit value below which a suppressed route can be used again.

Range: 1 through 20,000

Default: 750 (unitless)

suppress *number*—Cutoff (suppression) threshold. ***number*** is the figure-of-merit value above which a route is suppressed for use or inclusion in advertisements.

Range: 1 through 20,000

Default: 3000 (unitless)

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring BGP Flap Damping Parameters

dynamic-db

Syntax
dynamic-db;

Hierarchy Level
[edit logical-systems *logical-system-name* policy-options **as-path** *path-name*],
[edit logical-systems *logical-system-name* policy-options **as-path-group** *group-name*],
[edit logical-systems *logical-system-name* policy-options **community** *community-name*],
[edit logical-systems *logical-system-name* policy-options **condition** *condition-name*],
[edit logical-systems *logical-system-name* policy-options **policy-statement** *policy-statement-name*],
[edit logical-systems *logical-system-name* policy-options **prefix-list** *prefix-list-name*],
[edit policy-options **as-path** *path-name*],
[edit policy-options **as-path-group** *group-name*],
[edit policy-options **community** *community-name*],
[edit policy-options **condition** *condition-name*],
[edit policy-options **policy-statement** *policy-statement-name*],
[edit policy-options **prefix-list** *prefix-list-name*]

Release Information
Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Define routing policies and policy objects that reference policies configured in the dynamic database at the **[edit dynamic]** hierarchy level.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control-level—To add this statement to the configuration.

Related Documentation
- Configuring Routing Policies Based on Dynamic Database Configuration

family

Syntax

```
family family-name {
  filter filter-name {
    interface-specific;
    term term-name {
      from {
        match-conditions;
      }
      then {
        action;
        action-modifiers;
      }
    }
  }
}
```

Hierarchy Level [edit **firewall**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a firewall filter for IP version 4 or IP version 6.

Options ***family-name***—Version or type of addressing protocol:

- **any**—Filter packets based on protocol-independent match conditions.
- **ethernet-switching**—Filter Layer 2 (Ethernet) packets and Layer 3 (IP) packets.
- **inet**—Filter IPv4 packets.
- **inet6**—Filter IPv6 packets.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Firewall Filter Match Conditions and Actions for J-EX Series Switches on page 2728
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Configuring Firewall Filters (CLI Procedure) on page 2779
- Configuring Firewall Filters (J-Web Procedure) on page 2784
- Firewall Filters for J-EX Series Switches Overview on page 2721

filter

Syntax

```
filter filter-name {
  interface-specific;
  term term-name {
    from {
      match-conditions;
    }
    then {
      action;
      action-modifiers;
    }
  }
}
```

Hierarchy Level [edit **firewall family** *family-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure firewall filters.

Options

filter-name—Name that identifies the filter. The name can contain letters, numbers, and hyphens (-), and can be up to 64 characters long. To include spaces in the name, enclose it in quotation marks.

The remaining statements are explained separately.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- Firewall Filter Match Conditions and Actions for J-EX Series Switches on page 2728
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Configuring Firewall Filters (CLI Procedure) on page 2779
- Configuring Firewall Filters (J-Web Procedure) on page 2784
- Firewall Filters for J-EX Series Switches Overview on page 2721

filter

Syntax filter (input | output) *filter-name*;

Hierarchy Level [edit **interfaces** *interface-name* **unit** *logical-unit-number* **family** ethernet-switching],
[edit **interfaces** *interface-name* **unit** *logical-unit-number* **family** inet],
[edit **interfaces** *interface-name* **unit** *logical-unit-number* **family** inet6]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply a firewall filter to traffic entering the port or Layer 3 interface or exiting the Layer 3 interface.

Default All incoming traffic is accepted unmodified on the port or Layer 3 interface, and all outgoing traffic is sent unmodified from the port or Layer 3 interface.

Options ***filter-name***—Name of a firewall filter defined in the **filter** statement.

- **input**—Apply a firewall filter to traffic entering the port or Layer 3 interface.
- **output**—Apply a firewall filter to traffic exiting the Layer 3 interface.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Configuring Gigabit Ethernet Interfaces (CLI Procedure) on page 919
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 909
- Configuring Firewall Filters (CLI Procedure) on page 2779
- Configuring Firewall Filters (J-Web Procedure) on page 2784
- Firewall Filters for J-EX Series Switches Overview on page 2721
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

filter

Syntax filter (input | output) *filter-name*;

Hierarchy Level [edit **vlans** *vlan-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply a firewall filter to traffic coming into or exiting from the VLAN.

Default All incoming traffic is accepted unmodified to the VLAN, and all outgoing traffic is sent unmodified from the VLAN.

Options ***filter-name*** —Name of a firewall filter defined in a **filter** statement.

- **input**—Apply a firewall filter to VLAN ingress traffic.
- **output**—Apply a firewall filter to VLAN egress traffic.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Configuring Firewall Filters (CLI Procedure) on page 2779
- Configuring Firewall Filters (J-Web Procedure) on page 2784
- Firewall Filters for J-EX Series Switches Overview on page 2721

filter-specific

Syntax filter-specific;

Hierarchy Level [edit **firewall policer** *policer-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a policer to act as a filter-specific policer.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration

Related Documentation

- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Configuring Policers to Control Traffic Rates (CLI Procedure) on page 2788
- Understanding the Use of Policers in Firewall Filters on page 2752

firewall

Syntax

```
firewall {
  familyfamily-name {
    filter filter-name {
      interface-specific;
      term term-name {
        from {
          match-conditions;
        }
        then {
          action;
          action-modifiers;
        }
      }
    }
  }
  policerpolicer-name {
    filter-specific;
    if-exceeding {
      bandwidth-limit bps;
      burst-size-limit bytes;
    }
    then {
      policer-action;
    }
  }
}
```

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure firewall filters and policers.

The remaining statements are explained separately.

Required Privilege Level firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- Firewall Filter Match Conditions and Actions for J-EX Series Switches on page 2728
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Configuring Firewall Filters (CLI Procedure) on page 2779
- Configuring Policers to Control Traffic Rates (CLI Procedure) on page 2788
- Firewall Filters for J-EX Series Switches Overview on page 2721

from

Syntax

```
from {
  match-conditions;
}
```

Hierarchy Level [edit **firewall family** *family-name* **filter** *filter-name* **term** *term-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Match packet fields to values specified in a match condition. If the **from** statement is not included in a firewall filter configuration, all packets are considered to match and the actions and action modifiers in the **then** statement are taken.

Options

match-conditions —Conditions that define the values or fields that the incoming or outgoing packets must contain for a match. You can specify one or more match conditions. If you specify more than one, they all must match for a match to occur and for the action in the **then** statement to be taken.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- Firewall Filter Match Conditions and Actions for J-EX Series Switches on page 2728
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Configuring Firewall Filters (CLI Procedure) on page 2779
- Configuring Firewall Filters (J-Web Procedure) on page 2784
- Understanding Firewall Filter Match Conditions on page 2748

if-exceeding

Syntax

```
if-exceeding {
  bandwidth-limit bps;
  bandwidth-percentpercent
  burst-size-limit bytes;
}
```

Hierarchy Level

[edit **firewall policer** *policer-name*]
[edit logical-systems logical-system-name firewall policer policer-name]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure policer rate limits.

The **bandwidth-percent** statement is supported on routers only.

The remaining statements are explained separately.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Configuring Policers to Control Traffic Rates (CLI Procedure) on page 2788
- Understanding the Use of Policers in Firewall Filters on page 2752
- Rate Limiting
- Single-Rate Two-Color Policer Overview
- Configuring a Single-Rate Two-Color Policer

interface-specific

Syntax interface-specific;

Hierarchy Level [edit **firewall family** *family-name* **filter** *filter-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure interface-specific names for firewall counters.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Firewall Filter Match Conditions and Actions for J-EX Series Switches on page 2728
- Configuring Firewall Filters (CLI Procedure) on page 2779
- Configuring Firewall Filters (J-Web Procedure) on page 2784
- Firewall Filters for J-EX Series Switches Overview on page 2721

policer

Syntax

```
policer policer-name {
  filter-specific;
  if-exceeding {
    bandwidth-limit bps;
    bandwidth-percent percent
    burst-size-limit bytes;
  }
  then {
    policer-action;
  }
}
```

Hierarchy Level

[edit firewall]
[edit logical-systems *logical-system-name* firewall]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure policer rate limits and actions. To activate a policer, you must include the policer action modifier in the **then** statement in a firewall filter term. Each policer that you configure includes an implicit counter. To ensure term-specific packet counts, you configure a policer for each term in the filter that requires policing.

Options

policer-name—Name that identifies the policer. The name can contain letters, numbers, hyphens (-), and can be up to 64 characters long.

The remaining statements are explained separately.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Example: Combining CoS with MPLS on J-EX Series Switches on page 2883
- Configuring Policers to Control Traffic Rates (CLI Procedure) on page 2788
- Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect (CLI Procedure) on page 3111
- Configuring MPLS on Provider Edge Switches Using IP Over MPLS (CLI Procedure) on page 3107
- Configuring Firewall Filters (CLI Procedure) on page 2779
- Configuring Firewall Filters (J-Web Procedure) on page 2784
- Understanding the Use of Policers in Firewall Filters on page 2752
- Single-Rate Two-Color Policer Overview
- Configuring a Single-Rate Two-Color Policer

policy-statement

Syntax

```
policy-statement policy-name {
  term term-name {
    from {
      family family-name;
      match-conditions;
      policy subroutine-policy-name;
      prefix-list prefix-list-name;
      prefix-list-filter prefix-list-name match-type <actions>;
      route-filter destination-prefix match-type <actions>;
      source-address-filter source-prefix match-type <actions>;
    }
    to {
      match-conditions;
      policy subroutine-policy-name;
    }
    then actions;
  }
}
```

Hierarchy Level

[edit dynamic policy-options],
[edit logical-systems *logical-system-name* policy-options],
[edit policy-options]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Define a routing policy, including subroutine policies.

Options

actions—(Optional) One or more actions to take if the conditions match. The actions are described in Configuring Flow Control Actions.

family ***family-name***—(Optional) Specify an address family protocol. Specify **inet** for IPv4. Specify **inet6** for 128-bit IPv6, and to enable interpretation of IPv6 router filter addresses. For IS-IS traffic, specify **iso**. For IPv4 multicast VPN traffic, specify **inet-mvpn**. For IPv6 multicast VPN traffic, specify **inet6-mvpn**. For multicast-distribution-tree (MDT) IPv4 traffic, specify **inet-mdt**.

NOTE: When family is not specified, the routing device uses the default IPv4 setting.

from—(Optional) Match a route based on its source address.

match-conditions—(Optional in **from** statement; required in **to** statement) One or more conditions to use to make a match. The qualifiers are described in Configuring Match Conditions in Routing Policy Terms.

policy ***subroutine-policy-name***—Use another policy as a match condition within this policy. The name identifying the subroutine policy can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it

in quotation marks (“ ”). For information about how to configure subroutines, see Configuring Subroutines in Routing Policy Match Conditions.

policy-name—Name that identifies the policy. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose it in quotation marks (“ ”).

prefix-list ***prefix-list-name*** —Name of a list of IPv4 or IPv6 prefixes.

prefix-list-filter ***prefix-list-name***—Name of a prefix list to evaluate using qualifiers; ***match-type*** is the type of match (see Configuring Prefix List Filters), and ***actions*** is the action to take if the prefixes match.

route-filter ***destination-prefix match-type*** **<*actions*>**—(Optional) List of routes on which to perform an immediate match; ***destination-prefix*** is the IPv4 or IPv6 route prefix to match, ***match-type*** is the type of match (see Configuring Route Lists), and ***actions*** is the action to take if the ***destination-prefix*** matches.

source-address-filter ***source-prefix match-type*** **<*actions*>**—(Optional) Unicast source addresses in multiprotocol BGP (MBGP) and Multicast Source Discovery Protocol (MSDP) environments on which to perform an immediate match. ***source-prefix*** is the IPv4 or IPv6 route prefix to match, ***match-type*** is the type of match (see Configuring Route Lists), and ***actions*** is the action to take if the ***source-prefix*** matches.

term ***term-name***—Name that identifies the term.

to—(Optional) Match a route based on its destination address or the protocols into which the route is being advertised.

then—(Optional) Actions to take on matching routes. The actions are described in Configuring Flow Control Actions and Configuring Actions That Manipulate Route Characteristics.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Defining Routing Policies
- Configuring Routing Policies and Policy Objects in the Dynamic Database
- **dynamic-db on page 2817**

prefix-list

Syntax

```
prefix-list name {
  ip-addresses;
  apply-path path;
}
```

Hierarchy Level

[edit dynamic policy-options],
[edit logical-systems *logical-system-name* policy-options],
[edit policy-options]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.
Support for the **vpls** protocol family introduced in Junos OS Release 10.2 for J-EX Series switches.

Description

Define a list of IPv4 or IPv6 address prefixes for use in a routing policy statement or firewall filter statement.

Options

name—Name that identifies the list of IPv4or IPv6 address prefixes.

ip-addresses—List of IPv4 or IPv6 address prefixes, one IP address per line in the configuration.

The remaining statement is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Prefix Lists for Use in Routing Policy Match Conditions
- Configuring Routing Policies and Policy Objects in the Dynamic Database
- **dynamic-db on page 2817**

routing-instance

Syntax routing-instance *routing-instance-name*;

Hierarchy Level [edit **firewall** family inet **filter** *filter-name* **term** *term-name* **then**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify a specific virtual routing instance to which the switch sends matched packets.

Options ***routing-instance-name*** —Name of a virtual routing instance.

Required Privilege Level firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Filter-Based Forwarding on J-EX Series Switches on page 2773
- Configuring Virtual Routing Instances (CLI Procedure) on page 1142
- Understanding Filter-Based Forwarding for J-EX Series Switches on page 2753

term

Syntax

```
term term-name {
  from {
    match-conditions;
  }
  then {
    action;
    action-modifiers;
  }
}
```

Hierarchy Level [edit **firewall family** *family-name* **filter** *filter-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define a firewall filter term.

Options ***term-name*** —Name that identifies the term. The name can contain letters, numbers, and hyphens (-), and can be up to 64 characters long. To include spaces in the name, enclose it in quotation marks.

The remaining statements are explained separately.

Required Privilege Level firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- Firewall Filter Match Conditions and Actions for J-EX Series Switches on page 2728
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Configuring Firewall Filters (CLI Procedure) on page 2779
- Configuring Firewall Filters (J-Web Procedure) on page 2784
- Firewall Filters for J-EX Series Switches Overview on page 2721

then

Syntax

```
then {
  action;
  action-modifiers;
}
```

Hierarchy Level [edit **firewall family** *family-name* **filter** *filter-name* **term** *term-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a filter action.

Options

action—Actions to accept, discard, or forward packets that match all match conditions specified in a filter term.

action-modifiers—Additional actions to analyze, classify, count, or police packets that match all conditions specified in a filter term.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall-control—To add this statement to the configuration.

Related Documentation

- Firewall Filter Match Conditions and Actions for J-EX Series Switches on page 2728
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Example: Using Filter-Based Forwarding to Route Application Traffic to a Security Device on J-EX Series Switches on page 2773
- Configuring Firewall Filters (CLI Procedure) on page 2779
- Configuring Firewall Filters (J-Web Procedure) on page 2784
- Understanding Firewall Filter Match Conditions on page 2748

then

Syntax

```
then {
  policer-action;
}
```

Hierarchy Level

[edit **firewall policer** *policer-name*]
[edit logical-systems *logical-system-name* firewall policer *policer-name*]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Configure a policer action.

Options

policer-action—Actions to take are:

- **discard**—Discard traffic that exceeds the rate limits defined by the policer.
- **forwarding-class *class-name***—For routers only, classify traffic that exceeds the rate limits defined by the policer.
- **loss-priority**—Set the loss priority for traffic that exceeds the rate limits defined by the policer.

Required Privilege Level

firewall—To view this statement in the configuration.
firewall -control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Configuring Policers to Control Traffic Rates (CLI Procedure) on page 2788
- Configuring Firewall Filters (CLI Procedure) on page 2779
- Configuring Firewall Filters (J-Web Procedure) on page 2784
- Understanding the Use of Policers in Firewall Filters on page 2752
- Example: Configuring CoS for a PBB Network on MX Series Routers
- Single-Rate Two-Color Policer Overview
- Configuring a Single-Rate Two-Color Policer

CHAPTER 106

Operational Mode Commands for Firewall Filters

clear firewall

Syntax clear firewall (all | counter *counter-name* | filter *filter-name* | logical-system *logical-system-name*)

Syntax (J-EX Series Switch) clear firewall (all | counter *counter-name* | filter *filter-name*)

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Clear statistics about configured firewall filters.

NOTE: The clear firewall command cannot be used to clear the Routing Engine filter counters on a backup Routing Engine that is enabled for GRES.

Options

all—Clear the packet and byte counts for all filters.

counter *counter-name*—Clear the packet and byte counts for a filter counter that has been configured with the counter firewall filter action.

filter *filter-name*—Clear the packet and byte counts for the specified firewall filter.

logical-system *logical-system-name*—Clear the packet and byte counts for the specified logical system.

Required Privilege Level clear

Related Documentation

- **show firewall on page 2838**

List of Sample Output **clear firewall all on page 2836**

Output Fields When you enter this command, you are provided feedback on the status of your request.

clear firewall all

```
user@host> clear firewall all
```

clear firewall

Syntax

clear firewall
<all>
<counter *counter-name*>
<filter *filter-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Clear statistics about configured firewall filters.

Options

none—Clear the packet and byte counts for all firewall filter counters and clear the packet counts for all policer counters.

all—(Optional) Clear the packet and byte counts for all firewall filter counters and clear the packet counts for all policer counters.

counter *counter-name* —(Optional) Clear the packet and byte counts for the specified firewall filter counter.

filter *filter-name* —(Optional) Clear the packet and byte counts for the specified firewall filter.

Required Privilege Level

clear

Related Documentation

- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Verifying That Firewall Filters Are Operational on page 2799
- Verifying That Policers Are Operational on page 2800
- Firewall Filters for J-EX Series Switches Overview on page 2721
- Understanding the Use of Policers in Firewall Filters on page 2752

clear firewall (all)

```
user@host> clear firewall all
```

clear firewall (counter counter-name)

```
user@host> clear firewall counter port-filter-counter
```

clear firewall (filter filter-name)

```
user@host> clear firewall filter ingress-port-filter
```

show firewall

Syntax

show firewall
<filter *filter-name*>
<counter *counter-name*>
<log>
<logical-system (all | *logical-system-name*)>
<terse>

Syntax (J-EX Series Switch)

show firewall
<filter *filter-name*>
<counter *counter-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display statistics about configured firewall filters.

Options

none—(Optional) Display statistics about configured firewall filters.

filter *filter-name*—(Optional) Name of a configured filter.

counter *counter-name*—(Optional) Name of a filter counter.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular system.

log—(Optional) Display log entries for firewall filters.

terse—(Optional) Display firewall filter names only.

Required Privilege Level

view

Related Documentation

- **clear firewall on page 2836**

List of Sample Output

show firewall filter on page 2840
show firewall filter (Dynamic Input Filter) on page 2840
show firewall (Logical Systems) on page 2840

Output Fields

Table 350 on page 2839 lists the output fields for the **show firewall** command. Output fields are listed in the approximate order in which they appear.

Table 350: show firewall Output Fields

Field Name	Field Description
Filter	Name of a filter that has been configured with the **filter** statement at the **[edit firewall]** hierarchy level. When an interface-specific filter is displayed, the name of the filter is followed by the full interface name and by either **-i** for an input filter or **-o** for an output filter. When dynamic filters are displayed, the name of the filter is followed by the full interface name and by either **-in** for an input filter or **-out** for an output filter. When a logical system–specific filter is displayed, the name of the filter is prefixed with two underscore (__) characters and the name of the logical system (for example, **__ls1/filter1**).
Counters	Display filter counter information: • **Name**—Name of a filter counter that has been configured with the **counter** firewall filter action. • **Bytes**—Number of bytes that match the filter term under which the **counter** action is specified. • **Packets**—Number of packets that matched the filter term under which the **counter** action is specified.
Policers	Display policer information: • **Name**—Name of policer. • **Packets**—Number of packets that matched the filter term under which the policer action is specified. This is only the number of out-of-specification (out-of-spec) packet counts, not all packets policed by the policer.

show firewall filter

```
user@host> show firewall filter test
Filter: test
Counters:
Name                                Bytes              Packets
Counter-1                               0                    0
Counter-2                               0                    0
Policers:
Name                                                   Packets
Policer-1                                                    0
```

show firewall filter (Dynamic Input Filter)

```
user@host> show firewall filter dfwd-ge-5/0/0.1-in
Filter: dfwd-ge-5/0/0.1-in
Counters:
Name                                                  Bytes              Packets
c1-ge-5/0/0.1-in                                          0                    0
```

show firewall (Logical Systems)

```
user@host>show firewall

Filter: __lr1/test
Counters:
Name                                                  Bytes              Packets
icmp                                                    420                    5
Filter: __default_bpdu_filter__
Filter: __lr1/inet_filter1
Counters:
Name                                                  Bytes              Packets
inet_tcp_count                                            0                    0
inet_udp_count                                            0                    0
Filter: __lr1/inet_filter2
Counters:
Name                                                  Bytes              Packets
inet_icmp_count                                           0                    0
inet_pim_count                                            0                    0
Filter: __lr2/inet_filter1
Counters:
Name                                                  Bytes              Packets
inet_tcp_count                                            0                    0
inet_udp_count                                            0                    0
```

show firewall

Syntax

show firewall
<counter *counter-name*>
<filter *filter-name*>
log (detail | interface *interface-name*)
terse

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display statistics about configured firewall filters.

Options

none—Display statistics about all configured firewall filters, counters, and policers.

counter *counter-name*—(Optional) Display statistics about a particular firewall filter counter.

filter *filter-name*—(Optional) Display statistics about a particular firewall filter.

log (detail | interface *interface-name*)—(Optional) Display detailed log entries of firewall activity or log information about a specific interface.

terse—(Optional) Display firewall filter names only.

Required Privilege Level

view

Related Documentation

- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Verifying That Firewall Filters Are Operational on page 2799
- Verifying That Policers Are Operational on page 2800
- Firewall Filters for J-EX Series Switches Overview on page 2721
- Understanding the Use of Policers in Firewall Filters on page 2752

List of Sample Output

show firewall on page 2842
show firewall (filter filter-name) on page 2842
show firewall (counter counter-name) on page 2842
show firewall log on page 2842

Output Fields

Table 351 on page 2841 lists the output fields for the **show firewall** command. Output fields are listed in the approximate order in which they appear.

Table 351: show firewall Output Fields

Field Name	Field Description	Level of Output
Filter	Name of the filter that is configured with the **filter** statement at the **[edit firewall]** hierarchy level.	All levels

Table 351: show firewall Output Fields *(continued)*

Field Name	Field Description	Level of Output
Counters	Display filter counter information: • Name—Name of a filter counter that has been configured with the counter firewall filter action • Bytes—Number of bytes that match the filter term where the counter action was specified. • Packets—Number of packets that matched the filter term where the counter action was specified.	All levels
Policers	Display policer information: • Name—Name of policer. • Packets—Number of packets that matched the filter term where the policer action was specified. This is the number of packets that exceed the rate limits that the policer specifies.	All levels

show firewall

```
user@host> show firewall
Filter: egress-vlan-filter
Counters:
Name                                                Bytes              Packets
employee-web-counter                                    0                    0
Filter: ingress-port-filter
Counters:
Name                                                Bytes              Packets
ingress-port-counter                                    0                    0
Filter: ingress-port-voip-class-filter
Counters:
Name                                                Bytes              Packets
icmp-counter                                            0                    0
Policers:
Name                                               Packets
icmp-connection-policer                                 0
tcp-connection-policer                                  0
```

show firewall (filter filter-name)

```
user@host> show firewall filter egress-vlan-filter
Filter: egress-vlan-filter
Counters:
Name                                                Bytes              Packets
employee-web-counter                                    0                    0
```

show firewall (counter counter-name)

```
user@host> show firewall counter icmp-counter
Filter: ingress-port-voip-class-filter
Counters:
Name                                                Bytes              Packets
icmp-counter                                            0                    0
```

show firewall log

```
user@host> show firewall log
Log :

Time      Filter    Action Interface     Protocol        Src Addr
          Dest Addr
08:00:53  pfe       R      ge-1/0/1.0    ICMP            192.168.3.5
        192.168.3.4
```

```
08:00:52  pfe       R      ge-1/0/1.0    ICMP           192.168.3.5
          192.168.3.4
08:00:51  pfe       R      ge-1/0/1.0    ICMP           192.168.3.5
          192.168.3.4
08:00:50  pfe       R      ge-1/0/1.0    ICMP           192.168.3.5
          192.168.3.4
08:00:49  pfe       R      ge-1/0/1.0    ICMP           192.168.3.5
          192.168.3.4
08:00:48  pfe       R      ge-1/0/1.0    ICMP           192.168.3.5
          192.168.3.4
08:00:47  pfe       R      ge-1/0/1.0    ICMP           192.168.3.5
          192.168.3.4
```

show firewall log

Syntax

show firewall log
<detail>
<interface *interface-name*>
<logical-system (*logical-system-name* | all)>

Syntax (J-EX Series Switch)

show firewall log
<detail>
<interface *interface-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display log information about firewall filters.

Options

none—Display log information about firewall filters.

detail—(Optional) Display detailed information.

interface *interface-name*—(Optional) Display log information about a specific interface.

logical-system (*logical-system-name* | all)—(Optional) Perform this operation on all logical systems or on a particular system.

Required Privilege Level

view

List of Sample Output

show firewall log on page 2845
show firewall log detail on page 2845

Output Fields

Table 352 on page 2844 lists the output fields for the **show firewall log** command. Output fields are listed in the approximate order in which they appear.

Table 352: show firewall log Output Fields

Field Name	Field Description
Time of Log	Time that the event occurred.
Filter	Name of a filter that has been configured with the **filter** statement at the **[edit firewall]** hierarchy level. • A hyphen (-) indicates that the packet was handled by the Packet Forwarding Engine. • A space (no hyphen) indicates the packet was handled by the Routing Engine. • The notation **pfe** indicates packets logged by the Packet Forwarding Engine hardware filters.

Table 352: show firewall log Output Fields *(continued)*

Field Name	Field Description
Filter Action	Filter action: • **A**—Accept • **D**—Discard • **R**—Reject
Name of Interface	Ingress interface for the packet.
Name of protocol	Packet's protocol name: **egp**, **gre**, **icmp**, **ipip**, **ospf**, **pim**, **rsvp**, **tcp**, or **udp**.
Packet length	Length of the packet.
Source address	Packet's source address.
Destination address	Packet's destination address and port.

show firewall log

```
user@host>show firewall log
Time      Filter    Action Interface     Protocol  Src Addr      Dest Addr

13:10:12  pfe       D      rlsq0.902     ICMP      180.1.177.2   180.1.177.1

13:10:11  pfe       D      rlsq0.902     ICMP      180.1.177.2   180.1.177.1
```

show firewall log detail

```
user@host> show firewall log detail
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0Name of protocol: TCP, Packet Length: 50824, Source address:
172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 1020, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
Time of Log: 2004-10-13 10:37:17 PDT, Filter: f, Filter action: accept, Name of
interface: fxp0.0
Name of protocol: TCP, Packet Length: 49245, Source address: 172.17.22.108:829,
Destination address: 192.168.70.66:513
....
```

show interfaces filters

Syntax **show interfaces filters**
<interface-name>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display firewall filters that are configured on each interface in a system.

Options none—Display firewall filter information about all interfaces.

interface-name—(Optional) Display firewall filter information about a particular interface.

Required Privilege Level view

Related Documentation
- **show interfaces policers on page 2848**
- **show firewall on page 2841**

List of Sample Output **show interfaces filters on page 2846**
show interfaces filters <interface-name> on page 2847

Output Fields Table 353 on page 2846 lists the output fields for the **show interfaces filters** command. Output fields are listed in the approximate order in which they appear.

Table 353: show interfaces filters Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the physical interface.	All levels
Admin	Interface state: up or down.	All levels
Link	Link state: up or down.	All levels
Proto	Protocol that is configured on the interface.	All levels
Input Filter	Name of the firewall filter to be evaluated when packers are received on the interface.	All levels
Output Filter	Name of the firewall filter to be evaluated when packets are transmitted on the interface.	All levels

show interfaces filters

```
user@host> show interfaces filters
Interface       Admin Link Proto Input Filter        Output Filter
ge-0/0/0        up    down
ge-0/0/0.0      up    down eth-switch unknown
ge-0/0/1        up    down
ge-0/0/1.0      up    down eth-switch unknown
ge-0/0/2        up    down
ge-0/0/3        up    down
```

```
ge-0/0/4         up    down
ge-0/0/5         up    down
ge-0/0/6         up    down
ge-0/0/7         up    down
ge-0/0/8         up    down
ge-0/0/9         up    down
ge-0/0/10        up    down
ge-0/0/10.0      up    down
```

show interfaces filters <interface-name>

```
user@host> show interfaces filters ge-0/0/0
Interface       Admin Link Proto Input Filter         Output Filter
ge-0/0/0        up    down
ge-0/0/0.0      up    down eth-switch unknown
```

show interfaces policers

Syntax **show interfaces policers**
<interface-name>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display all policers that are configured on each interface in a system.

Options none—Display policer information about all interfaces.

interface-name—(Optional) display firewall filters information about a particular interface.

Required Privilege Level view

Related Documentation
- **show interfaces filters on page 2846**
- **show policer on page 2850**

List of Sample Output
show interfaces policers on page 2848
show interfaces policers on page 2849
show interfaces policers (interface-name) on page 2849

Output Fields Table 354 on page 2848 lists the output fields for the **show interfaces policers** command. Output fields are listed in the approximate order in which they appear.

Table 354: show interfaces policers Output Fields

Field Name	Field Description	Level of Output
Interface	Name of the interface.	All levels
Admin	Interface state: up or down.	All levels
Link	Link state: up or down.	All levels
Proto	Protocol configured on the interface.	All levels
Input Policer	Policer to be evaluated when packets are received on the interface. It has the format ***interface-name*-in-policer**.	All levels
Output Policer	Policer to be evaluated when packets are transmitted on the interface. It has the format ***interface-name*-out-policer**.	All levels

show interfaces policers

```
user@host> show interfaces policers
Interface       Admin Link Proto Input Policer        Output Policer
ge-0/0/0        up    down
ge-0/0/0.0      up    down
                            eth-switch

Interface       Admin Link Proto Input Policer        Output Policer
```

```
ge-0/0/1        up    down
ge-0/0/1.0      up    down
                            eth-switch
Interface       Admin Link Proto Input Policer         Output Policer
ge-0/0/2        up    down
ge-0/0/3        up    down
ge-0/0/4        up    down
ge-0/0/5        up    down
ge-0/0/6        up    down
ge-0/0/7        up    down
ge-0/0/8        up    down
ge-0/0/9        up    down
ge-0/0/10       up    down
ge-0/0/10.0     up    down
                            eth-switch
```

show interfaces policers

```
user@host> show interfaces policers
Interface       Admin Link Proto Input Policer         Output Policer
ge-0/0/0        up    down
ge-0/0/0.0      up    down
                            eth-switch

Interface       Admin Link Proto Input Policer         Output Policer
ge-0/0/1        up    down
ge-0/0/1.0      up    down
                            eth-switch
Interface       Admin Link Proto Input Policer         Output Policer
ge-0/0/2        up    down
ge-0/0/3        up    down
ge-0/0/4        up    down
ge-0/0/5        up    down
ge-0/0/6        up    down
ge-0/0/7        up    down
ge-0/0/8        up    down
ge-0/0/9        up    down
ge-0/0/10       up    down
ge-0/0/10.0     up    down
                            eth-switch
```

show interfaces policers (interface-name)

```
user@host> show interfaces policers ge-0/0/1
Interface       Admin Link Proto Input Policer         Output Policer
ge-0/0/0        up    down
ge-0/0/0.0      up    down
                            eth-switch
```

show policer

Syntax **show policer**
<policer-name>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display statistics about configured policers.

Options none—Display the count of policed packets for all configured policers in the system.

policer-name—(Optional) Display the count of policed packets for the specified policer.

Required Privilege Level view

Related Documentation

- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Verifying That Firewall Filters Are Operational on page 2799
- Verifying That Policers Are Operational on page 2800
- Firewall Filters for J-EX Series Switches Overview on page 2721
- Understanding the Use of Policers in Firewall Filters on page 2752

List of Sample Output **show policer on page 2850**
show policer (policer-name) on page 2851

Output Fields Table 355 on page 2850 lists the output fields for the **show policer** command. Output fields are listed in the approximate order in which they appear.

Table 355: show policer Output Fields

Field Name	Field Description	Level of Output
Filter	Name of filter that is configured with the **filter** statement at the **[edit firewall]** hierarchy level.	All levels
Policers	Display policer information: • Filter—Name of filter that specifies the policer action. • Name—Name of policer. • Packets—Number of packets that matched the filter term where the policer action is specified. This is the number of packets that exceed the rate limits that the policer specifies.	All levels

show policer

```
user@host> show policer
Filter: egress-vlan-filter
Filter: ingress-port-filter
Policers:
Name                                                Packets
```

```
icmp-connection-policer                                  0
tcp-connection-policer                                   0
Filter: ingress-vlan-rogue-block
```

show policer (policer-name)

```
user@host> show policer tcp-connection-policer
Filter: ingress-port-filter
Policers:
Name                                                Packets
tcp-connection-policer                                    0
```

show policy

Syntax

show policy
<logical-system (all | *logical-system-name*)>
<*policy-name*>

Syntax (J-EX Series Switch)

show policy
<*policy-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display information about configured routing policies.

Options

none—List the names of all configured routing policies.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

policy-name—(Optional) Show the contents of the specified policy.

Required Privilege Level

view

Related Documentation

- **show policy damping on page 1886**

List of Sample Output

show policy on page 2852
show policy policy-name on page 2853
show policy (Multicast Scoping) on page 2853

Output Fields

Table 356 on page 2852 lists the output fields for the **show policy** command. Output fields are listed in the approximate order in which they appear.

Table 356: show policy Output Fields

Field Name	Field Description
policy-name	Name of the policy listed.
term	Policy term listed.
from	Match condition for the policy.
then	Action for the policy.

show policy

```
user@host> show policy
Configured policies:
__vrf-export-red-internal__
__vrf-import-red-internal__
```

```
red-export
all_routes
```

show policy policy-name

```
user@host> show policy test-statics
Policy test-statics:
    from
         3.0.0.0/8  accept
         3.1.0.0/16  accept
    then reject
```

show policy (Multicast Scoping)

```
user@host> show policy test-statics
Policy test-statics:
    from
         multicast-scoping == 8
```

show policy conditions

Syntax

show policy conditions
<*condition-name*>
<detail>
<dynamic>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)

show policy conditions
<*condition-name*>
<detail>
<dynamic>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display all the configured conditions as well as the routing tables with which the configuration manager is interacting. If the **detail** keyword is included, the output also displays dependent routes for each condition.

Options

none—Display all configured conditions and associated routing tables.

condition-name—(Optional) Display information about the specified condition only.

detail—(Optional) Display the specified level of output.

dynamic—(Optional) Display information about the conditions in the dynamic database.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show policy conditions detail on page 2855

Output Fields

Table 357 on page 2854 lists the output fields for the **show policy conditions** command. Output fields are listed in the approximate order in which they appear.

Table 357: show policy conditions Output Fields

Field Name	Field Description	Level of Output
Condition	Name of configured condition.	All levels
event	Condition type. If the **if-route-exists** option is configured, the event type is: **Existence of a route in a specific routing table**.	All levels
Dependent routes	List of routes dependent on the condition, along with the latest generation number.	**detail**
Condition tables	List of routing tables associated with the condition, along with the latest generation number and number of dependencies.	All levels

Table 357: show policy conditions Output Fields *(continued)*

Field Name	Field Description	Level of Output
If-route-exists conditions	List of conditions configured to look for a route in the specified table.	All levels

show policy conditions detail

```
user@host> show policy conditions detail
Configured conditions:
Condition cond1, event: Existence of a route in a specific routing table
Dependent routes:
 4.4.4.4/32, generation 3
 6.6.6.6/32, generation 3
 10.10.10.10/32, generation 3

Condition cond2, event: Existence of a route in a specific routing table
Dependent routes:
None

Condition tables:
Table inet.0, generation 4, dependencies 3, If-route-exists conditions: cond1
cond2
```

test policy

Syntax test policy *policy-name prefix*

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Test a policy configuration to determine which prefixes match routes in the routing table.

Options *policy-name*—Name of a policy.

prefix—Destination prefix to match.

Additional Information All prefixes in the default unicast routing table (**inet.0**) that match prefixes that are the same as or longer than the specific prefix are processed by the **from** clause in the specified policy. All prefixes accepted by the policy are displayed. The **test policy** command evaluates a policy differently from the Border Gateway Protocol (BGP) import process. When testing a policy that contains an **interface** match condition in the **from** clause, the **test policy** command uses the match condition. In contrast, BGP does not use the **interface** match condition when evaluating the policy against routes learned from internal BGP (IBGP) or external BGP (EGBP) multihop peers.

Required Privilege Level view

Related Documentation
- **show policy damping on page 1886**

List of Sample Output **test policy on page 2856**

Output Fields For information about output fields, see the output field tables for the **show route** command, the **show route detail** command, the **show route extensive** command, or the **show route terse** command.

test policy

```
user@host> test policy test-statics 3.0.0.1/8
inet.0: 44 destinations, 44 routes (44 active, 0 holddown, 0 hidden)
Prefixes passing policy:

3.0.0.0/8           *[BGP/170] 16:22:46, localpref 100, from 10.255.255.41
                    AS Path: 50888 I
                     > to 10.11.4.32 via en0.2, label-switched-path l2
3.3.3.1/32          *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
                     > to 10.0.4.7 via fxp0.0
3.3.3.2/32          *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
                     > to 10.0.4.7 via fxp0.0
3.3.3.3/32          *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
                     > to 10.0.4.7 via fxp0.0
3.3.3.4/32          *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
                     > to 10.0.4.7 via fxp0.0
Policy test-statics: 5 prefixes accepted, 0 prefixes rejected
```

PART 21

Class of Service

CHAPTER 107

Class of Service (CoS)—Overview

Junos OS CoS for J-EX Series Switches Overview

When a network experiences congestion and delay, some packets must be dropped. Junos OS class of service (CoS) divides traffic into classes to which you can apply different levels of throughput and packet loss when congestion occurs. This allows packet loss to happen according to rules that you configure.

For interfaces that carry IPv4, IPv6, and MPLS traffic, you can configure Junos OS CoS features to provide multiple classes of service for different applications. CoS also allows you to rewrite the Differentiated Services code point (DSCP), IP precedence, 802.1p, or EXP CoS bits of packets egressing out of an interface, thus allowing you to tailor packets for the remote peers' network requirements. See "Understanding Using CoS with MPLS Networks on J-EX Series Switches" on page 2880 for more information about CoS for MPLS networks.

CoS provides multiple classes of service for different applications. You can configure multiple forwarding classes for transmitting packets, define which packets are placed into each output queue, and schedule the transmission service level for each queue.

In designing CoS applications, you must give careful consideration to your service needs and thoroughly plan and design your CoS configuration to ensure consistency and interoperability across all platforms in a CoS domain.

Because J-EX Series Switches implement CoS in hardware rather than in software, you can experiment with and deploy CoS features without affecting packet-forwarding and switching performance.

NOTE: CoS policies can be enabled or disabled on each interface of a J-EX Series switch. Also, each physical and logical interface on the switch can have custom CoS rules associated with it. When CoS is used in an MPLS network, there are some additional restrictions. See "Understanding Using CoS with MPLS Networks on J-EX Series Switches" on page 2880.

How Junos OS CoS Works

Junos OS CoS works by examining traffic entering at the edge of your network. The switches classify traffic into defined service groups to provide the special treatment of traffic across the network. For example, voice traffic can be sent across certain links, and data traffic can use other links. In addition, the data traffic streams can be serviced differently along the network path. As the traffic leaves the network at the far edge, you can rewrite the traffic to meet the policies of the targeted peer.

To support CoS, you must configure each switch in the network. Generally, each switch examines the packets that enter it to determine their CoS settings. These settings then dictate which packets are transmitted first to the next downstream switch. Switches at

the edges of the network might be required to alter the CoS settings of the packets that enter the network to classify the packets into the appropriate service groups.

Figure 79 on page 2861 represents the network scenario of an enterprise. Switch A is receiving traffic from various network nodes such as desktop computers, servers, surveillance cameras, and VoIP telephones. As each packet enters, Switch A examines the packet's CoS settings and classifies the traffic into one of the groupings defined by the enterprise. This definition allows Switch A to prioritize resources for servicing the traffic streams it receives. Switch A might alter the CoS settings of the packets to better match the enterprise's traffic groups.

When Switch B receives the packets, it examines the CoS settings, determines the appropriate traffic groups, and processes the packets according to those settings. It then transmits the packets to Switch C, which performs the same actions. Switch D also examines the packets and determines the appropriate groups. Because Switch D sits at the far end of the network, it can rewrite the CoS bits of the packets before transmitting them.

Figure 79: Packet Flow Across the Network

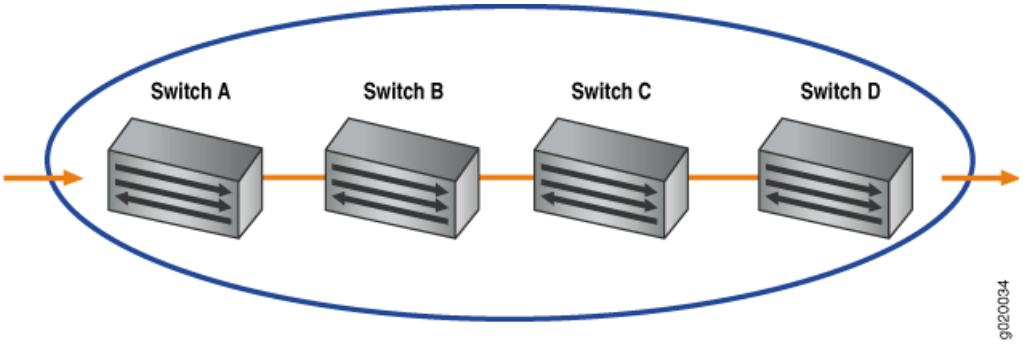

Default CoS Behavior on J-EX Series Switches

If you do not configure any CoS settings on the switch, the software performs some CoS functions to ensure that user traffic and protocol packets are forwarded with minimum delay when the network is experiencing congestion. Some CoS settings, such as classifiers, are automatically applied to each logical interface that you configure. Other settings, such as rewrite rules, are applied only if you explicitly associate them with an interface.

Related Documentation

- Understanding Junos OS CoS Components for J-EX Series Switches on page 2862
- Understanding Junos OS EZQoS for CoS Configurations on J-EX Series Switches on page 2879
- Example: Configuring CoS on J-EX Series Switches on page 2883
- Example: Combining CoS with MPLS on J-EX Series Switches on page 2898

Understanding Junos OS CoS Components for J-EX Series Switches

This topic describes the Junos OS class-of-service (CoS) components for J-EX Series Switches:

Code-Point Aliases

A code-point alias assigns a name to a pattern of code-point bits. You can use this name instead of the bit pattern when you configure other CoS components such as classifiers, drop-profile maps, and rewrite rules.

Policers

Policers limit traffic of a certain class to a specified bandwidth and *burst size*. Packets exceeding the policer limits can be discarded. You define policers with filters that can be associated with input interfaces.

For more information about policers, see “Understanding the Use of Policers in Firewall Filters” on page 2752.

NOTE: You can configure policers to discard packets that exceed the rate limits. If you want to configure CoS parameters such as loss-priority and forwarding-class, you must use firewall filters.

Classifiers

Packet classification associates incoming packets with a particular CoS servicing level. In the Junos OS, *classifiers* associate packets with a forwarding class and loss priority and assign packets to output queues based on the associated forwarding class. Junos OS supports two general types of classifiers:

- Behavior aggregate or CoS value traffic classifiers—Examines the CoS value in the packet header. The value in this single field determines the CoS settings applied to the packet. BA classifiers allow you to set the forwarding class and loss priority of a packet based on the Differentiated Services code point (DSCP) value, IP precedence value, and IEEE 802.1p value.
- Multifield traffic classifiers—Examines multiple fields in the packet such as source and destination addresses and source and destination port numbers of the packet. With

multifield classifiers, you set the forwarding class and loss priority of a packet based on firewall filter rules.

Forwarding Classes

Forwarding classes group the packets for transmission. Based on forwarding classes, you assign packets to output queues. Forwarding classes affect the forwarding, scheduling, and marking policies applied to packets as they transit a switch. By default, four categories of forwarding classes are defined: best effort, assured forwarding, expedited forwarding, and network control. For J-EX Series switches, 16 forwarding classes are supported, providing granular classification capability.

Tail Drop Profiles

Drop profile is a mechanism that defines parameters that allow packets to be dropped from the network. Drop profiles define the meanings of the loss priorities. When you configure drop profiles you are essentially setting the value for queue fullness. The queue fullness represents a percentage of the queue used to store packets in relation to the total amount that has been allocated for that specific queue.

Loss priorities set the priority of dropping a packet. Loss priority affects the scheduling of a packet without affecting the packet's relative ordering. You can use the loss priority setting to identify packets that have experienced congestion. Typically you mark packets exceeding some service level with a high loss priority.

Schedulers

Each switch interface has multiple queues assigned to store packets. The switch determines which queue to service based on a particular method of scheduling. This process often involves determining which type of packet should be transmitted before another. You can define the priority, bandwidth, delay buffer size, and tail drop profiles to be applied to a particular queue for packet transmission.

A scheduler map associates a specified forwarding class with a scheduler configuration. You can associate up to four user-defined scheduler maps with the interfaces.

Rewrite Rules

A *rewrite rule* sets the appropriate CoS bits in the outgoing packet, thus allowing the next downstream device to classify the packet into the appropriate service group. Rewriting, or marking, outbound packets is useful when the switch is at the border of a network and must alter the CoS values to meet the policies of the targeted peer.

NOTE: Rewrite rules are applied when the packets are routed. Rewrite rules are not applied when the packets are forwarded.

Egress firewall filters can also assign forwarding class and loss priority so that the packets are rewritten based on forwarding class and loss priority.

Related Documentation

- Understanding CoS Code-Point Aliases on page 2864
- Understanding CoS Classifiers on page 2867
- Understanding CoS Forwarding Classes on page 2870
- Understanding CoS Tail Drop Profiles on page 2872
- Understanding CoS Schedulers on page 2873
- Understanding CoS Two-Color Marking on page 2876
- Understanding CoS Rewrite Rules on page 2876
- Example: Configuring CoS on J-EX Series Switches on page 2883

Understanding CoS Code-Point Aliases

A code-point alias assigns a name to a pattern of code-point bits. You can use this name instead of the bit pattern when you configure other CoS components such as classifiers, drop-profile maps, and rewrite rules.

Behavior aggregate classifiers use class-of-service (CoS) values such as Differentiated Services code points (DSCPs), IP precedence, and IEEE 802.1 bits to associate incoming packets with a particular CoS servicing level. On a switch, you can assign a meaningful name or alias to the CoS values and use this alias instead of bits when configuring CoS components. These aliases are not part of the specifications but are well known through usage. For example, the alias for DSCP 101110 is widely accepted as ef (expedited forwarding).

When you configure classes and define classifiers, you can refer to the markers by alias names. You can configure user-defined classifiers in terms of alias names. If the value of an alias changes, it alters the behavior of any classifier that references it.

You can configure code-point aliases for the following type of CoS markers:

- dscp—Handles incoming IPv4 packets.
- ieee-802.1—Handles Layer 2 CoS.
- inet-precedence—Handles incoming IPv4 packets. IP precedence mapping requires only the upper three bits of the DSCP field.

This topic covers:

- Default Code-Point Aliases on page 2864

Default Code-Point Aliases

Table 358 on page 2865 shows the default mappings between the bit values and standard aliases.

Table 358: Default Code-Point Aliases

CoS Value Types	Mapping
DSCP CoS Values	
ef	101110
af11	001010
af12	001100
af13	001110
af21	010010
af22	010100
af23	010110
af31	011010
af32	011100
af33	011110
af41	100010
af42	100100
af43	100110
be	000000
cs1	001000
cs2	010000
cs3	011000
cs4	100000
cs5	101000
nc1/cs6	110000
nc2/cs7	111000
IEEE 802.1p CoS Values	
be	000

Table 358: Default Code-Point Aliases *(continued)*

CoS Value Types	Mapping
be1	001
ef	100
ef1	101
af11	010
af12	011
nc1/cs6	110
nc2/cs7	111
Legacy IP Precedence CoS Values	
be	000
be1	001
ef	010
ef1	011
af11	100
af12	101
nc1/cs6	110
nc2/cs7	111

Related Documentation

- Understanding Junos OS CoS Components for J-EX Series Switches on page 2862
- Example: Configuring CoS on J-EX Series Switches on page 2883
- Defining CoS Code-Point Aliases (CLI Procedure) on page 2914
- Defining CoS Code-Point Aliases (J-Web Procedure) on page 2912

Understanding CoS Classifiers

Packet classification associates incoming packets with a particular class-of-service (CoS) servicing level. Classifiers associate packets with a forwarding class and loss priority and assign packets to output queues based on the associated forwarding class. There are two general types of classifiers:

- Behavior aggregate (BA) classifiers
- Multifield (MF) classifiers

You can configure both a BA classifier and an MF classifier on an interface. If you do this, the BA classification is performed first and then the MF classification. If the two classification results conflict, the MF classification result overrides the BA classification result.

NOTE: When a source media access control (MAC) address is learned, the frame that contains the source MAC address is always sent out on queue 0 while egressing from the network interface, irrespective of the classifier applied to the ingress interface.

On J-EX8200 Ethernet Switches, you can specify BA classifiers for bridged multidestination traffic and IP multidestination traffic. The BA classifier for multicast packets is applied to all interfaces on the J-EX8200 switch.

This topic describes:

- Behavior Aggregate Classifiers on page 2867
- Multifield Classifiers on page 2869

Behavior Aggregate Classifiers

The behavior aggregate classifier maps a CoS value to a forwarding class and loss priority. The forwarding class determines the output queue. The loss priority is used by a scheduler to control packet discard during periods of congestion.

There are three types of BA classifiers:

- Differentiated Services Code Point (DSCP) for IP DiffServ
- IP precedence bits
- IEEE 802.1p CoS bits

BA classifiers are based on fixed-length fields, which makes them computationally more efficient than MF classifiers. Therefore core devices, which handle high traffic volumes, are normally configured to perform BA classification.

Default Behavior Aggregate Classification

The Junos OS automatically assigns implicit default classifiers to all logical interfaces based on the type of interface. Table 359 on page 2868 lists different types of interfaces and the corresponding implicit default classifiers.

Table 359: Default BA Classification

Type of Interface	Default BA Classification
Trunk interface	**ieee8021p-default**
Layer 3 interface (IPv4)	**dscp-default**
Layer 3 interface (IPv6)	**dscp-ipv6-default**
Access interface	Untrusted
Routed VLAN interface (RVI)	No default classification

When you explicitly associate a classifier with a logical interface, you are in effect overriding the implicit default classifier with an explicit classifier.

On J-EX4200 switches, you can apply classifier rules for each interface. Table 360 on page 2868 describes the different classifier types you can configure on Layer 2 and Layer 3 interfaces.

Table 360: Allowed BA Classification

Type of Interface	Allowed BA Classification
Layer 2 interface	IEEE 802.1p, IP Precedence, DSCP, DSCP IPv6
Layer 3 interface (IPv4)	IEEE 802.1p, IP Precedence, DSCP
Layer 3 interface (IPv6)	IEEE 802.1p, IP Precedence, DSCP IPv6

You can configure all the allowed classifier types on the same logical interface or on different logical interfaces. If you need to apply all classifier rules on the same logical interface, configure the classifier rules allowed for both IPv4 and IPv6 on the logical interface.

If you have not explicitly associated a classifier with a logical interface, the default classifiers are assigned and classification works as follows:

- If the logical interface is configured with an IPv4 address, DSCP classifier is assigned by default, and IPv4 and IPv6 packets are classified using the DSCP classifier.
- If the logical interface is configured with an IPv6 address, DSCP IPv6 classifier is assigned by default, and IPv4 and IPv6 packets are classified using the DSCP IPv6 classifier.

NOTE: On J-EX8200 switches, only one classifier of type DSCP and of type IEEE 802.1p can be applied to an interface.

You can configure routed VLAN interfaces (RVIs) to classify packets. After you do this, the User Priority (UP) bits in the incoming packets are rewritten according to the default IEEE 802.1p rewrite rule, except on J-EX8200 switches. On J-EX8200 switches, you must explicitly assign the default IEEE 802.1p rewrite rule to RVIs.

NOTE: By default, all BA classifiers classify traffic into either the best-effort forwarding class or the network-control forwarding class.

Multifield Classifiers

Multifield classifiers examine multiple fields in a packet such as source and destination addresses and source and destination port numbers of the packet. With MF classifiers, you set the forwarding class and loss priority of a packet based on firewall filter rules.

MF classification is normally performed at the network edge because of the general lack of DSCP or IP precedence support in end-user applications. On an edge switch, an MF classifier provides the filtering functionality that scans through a variety of packet fields to determine the forwarding class for a packet. Typically, a classifier performs matching operations on the selected fields against a configured value.

Related Documentation

- Understanding Junos OS Components for J-EX Series Switches on page 2862
- Example: Configuring CoS on J-EX Series Switches on page 2883
- Defining CoS Classifiers (CLI Procedure) on page 2914
- Defining CoS Classifiers (J-Web Procedure) on page 2916

Understanding CoS Forwarding Classes

It is helpful to think of forwarding classes as output queues. In effect, the end result of classification is the identification of an output queue for a particular packet. For a classifier to assign an output queue to each packet, it must associate the packet with one of the following forwarding classes:

- expedited-forwarding (ef)—Provides a low loss, low latency, low jitter, assured bandwidth, end-to-end service.
- assured-forwarding (af)—Provides a group of values you can define and includes four subclasses: AF1, AF2, AF3, and AF4, each with two drop probabilities: low and high.
- best-effort (be)—Provides no service profile. Loss priority is typically not carried in a class-of-service (CoS) value.
- network-control (nc)—Supports protocol control and thus is typically high priority.
- multicast best-effort (mcast-be)—Used for high-priority multicast packets.
- multicast assured-forwarding (mcast-af)—Provides two drop profiles, high and low, for multicast packets.
- multicast best-effort (mcast-be)—Provides no service profile for multicast packets.

NOTE: The forwarding classes multicast expedited-forwarding, multicast assured-forwarding, and multicast best-effort are applicable only to J-EX8200 Ethernet Switches.

J-EX Series Switches support up to 16 forwarding classes, thus allowing granular packet classification. For example, you can configure multiple classes of EF traffic such as EF, EF1, and EF2.

J-EX Series switches support up to eight output queues. Therefore, if you configure more than eight forwarding classes, you must map multiple forwarding classes to single output queues.

- Default Forwarding Classes on page 2870

Default Forwarding Classes

Table 361 on page 2871 shows the four default forwarding classes defined for unicast traffic, and Table 362 on page 2871 shows the three default forwarding classes defined for multicast traffic.

NOTE: The default forwarding classes for multicast traffic are applicable only to J-EX8200 switches.

If desired, you can rename the forwarding classes associated with the queues supported on your switch. Assigning a new class name to an output queue does not alter the default

classification or scheduling that is applicable to that queue. CoS configurations can be quite complicated, so unless it is required by your scenario, we recommend that you not alter the default class names or queue number associations.

Table 361: Default Forwarding Classes for Unicast Packets

Forwarding Class Name	Comments
best-effort (be)	The software does not apply any special CoS handling to packets with 000000 in the DiffServ field. This is a backward compatibility feature. These packets are usually dropped under congested network conditions.
expedited-forwarding (ef)	The software delivers assured bandwidth, low loss, low delay, and low delay variation (jitter) end-to-end for packets in this service class. The software accepts excess traffic in this class, but in contrast to the assured forwarding class, the out-of-profile expedited-forwarding class packets can be forwarded out of sequence or dropped.
assured-forwarding (af)	The software offers a high level of assurance that the packets are delivered as long as the packet flow from the customer stays within a certain service profile that you define. The software accepts excess traffic, but it applies a tail drop profile to determine if the excess packets are dropped and not forwarded. Up to two drop probabilities (low and high) are defined for this service class.
network-control (nc)	The software delivers packets in this service class with a high priority. (These packets are not delay-sensitive.) Typically, these packets represent routing protocol hello or keep alive messages. Because loss of these packets jeopardizes proper network operation, packet delay is preferable to packet discard.

Table 362: Default Forwarding Classes for Multicast Packets

Forwarding Class Name	Comments
multicast best-effort (mcast-be)	The software does not apply any special CoS handling to the multicast packets. These packets are usually dropped under congested network conditions.
multicast expedited-forwarding (mcast-ef)	The software delivers assured bandwidth, low loss, low delay, and low delay variation (jitter) end-to-end for multicast packets in this service class. The software accepts excess traffic in this class, but in contrast to the multicast assured forwarding class, out-of-profile multicast expedited-forwarding class packets can be forwarded out of sequence or dropped.
multicast assured-forwarding (mcast-af)	The software offers a high level of assurance that the multicast packets are delivered as long as the packet flow from the customer stays within a certain service profile that you define. The software accepts excess traffic, but it applies a tail drop profile to determine if the excess packets are dropped and not forwarded. Up to two drop probabilities (low and high) are defined for this service class.

The following rules govern queue assignment:

- CoS configurations that specify more queues than the switch can support are not accepted. The commit fails with a detailed message that states the total number of queues available.
- All default CoS configurations are based on queue number. The name of the forwarding class that shows up when the default configuration is displayed is the forwarding class currently associated with that queue.

Related Documentation

- Understanding Junos OS CoS Components for J-EX Series Switches on page 2862
- Example: Configuring CoS on J-EX Series Switches on page 2883
- Defining CoS Forwarding Classes (CLI Procedure) on page 2918
- Defining CoS Forwarding Classes (J-Web Procedure) on page 2918

Understanding CoS Tail Drop Profiles

Tail drop profile is a congestion management mechanism that allows switch to drop arriving packets when queue buffers become full or begin to overflow.

Tail drop profiles define the meanings of the loss priorities. When you configure tail drop profiles you are essentially setting the value for queue fullness. The queue fullness represents a percentage of the memory used to store packets in relation to the total amount that has been allocated for that specific queue.

The queue fullness defines the delay-buffer bandwidth, which provides packet buffer space to absorb burst traffic up to the specified duration of delay. Once the specified delay buffer becomes full, packets with 100 percent drop probability are dropped from the tail of the buffer.

On J-EX Series Switches, drop probability is implicitly set to **100 percent** and it cannot be modified.

You specify drop probabilities in the drop profile section of the CoS configuration hierarchy and reference them in each scheduler configuration.

By default, if you do not configure any drop profile, tail drop profile is in effect and functions as the primary mechanism for managing congestion. In the default tail drop profile, when the fill level is 0 percent, the drop probability is 0 percent. When the fill level is 100 percent, the drop probability is 100 percent.

NOTE: The default drop profile associated with the packets whose loss priority is low cannot be modified. You can configure custom drop profile only for those packets whose loss priority is high.

Related Documentation

- Understanding Junos OS CoS Components for J-EX Series Switches on page 2862
- Example: Configuring CoS on J-EX Series Switches on page 2883
- Configuring CoS Tail Drop Profiles (CLI Procedure) on page 2925

Understanding CoS Schedulers

You use schedulers to define the properties of output queues. These properties include the amount of interface bandwidth assigned to the queue, the size of the memory buffer allocated for storing packets, the priority of the queue, and the drop profiles associated with the queue.

You associate the schedulers with forwarding classes by means of scheduler maps. You can then associate each scheduler map with an interface, thereby configuring the queues, packet schedulers, and tail drop processes that operate according to this mapping.

Default Schedulers

Each forwarding class has an associated scheduler priority. Only two forwarding classes, best-effort (queue0) and network-control (queue7) are used in the default configuration.

NOTE: On J-EX8200 Ethernet Switches three forwarding classes—best-effort (queue0), multicast best-effort (queue2), and network-control (queue7)—are used in the default configuration.

By default, the best-effort forwarding class (queue 0) receives 95 percent of the bandwidth and buffer space for the output link, and the network-control forwarding class (queue 7) receives 5 percent. The default drop profile causes the buffer to fill completely and then to discard all incoming packets until it has free space.

NOTE: On J-EX8200 switches, by default, the best-effort forwarding class (queue 0) receives 75 percent of the bandwidth, the multicast best-effort forwarding class (queue 2) receives 20 percent of the bandwidth and buffer space for the output link, and the network-control forwarding class (queue 7) receives 5 percent.

The expedited-forwarding and assured-forwarding classes have no scheduler because no resources are assigned to queue 5 and queue 1, by default. However, you can manually configure resources for the expedited-forwarding and assured-forwarding classes.

Also by default, each queue can exceed the assigned bandwidth if additional bandwidth is available from other queues. When a forwarding class does not fully use the allocated transmission bandwidth, the remaining bandwidth can be used by other forwarding

classes if they receive a larger amount of offered load than their allocated bandwidth allows.

Transmission Rate

The transmission-rate control determines the actual traffic bandwidth from each forwarding class you configure. The rate is specified in bits per second. Each queue is allocated some portion of the bandwidth of the outgoing interface.

This bandwidth amount can be a fixed value, such as 1 megabit per second (Mbps), a percentage of the total available bandwidth, or the rest of the available bandwidth. You can allow transmission bandwidth to exceed the configured rate if additional bandwidth is available from other queues. In case of congestion, configured amount of transmission rate is guaranteed for the queue. This property allows you to ensure that each queue receives the amount of bandwidth appropriate to its level of service.

Scheduler Buffer Size

To control congestion at the output stage, you can configure the delay-buffer bandwidth. The delay-buffer bandwidth provides packet buffer space to absorb burst traffic up to the specified duration of delay. Once the specified delay buffer becomes full, packets with 100 percent drop probability are dropped from the tail of the buffer.

The default scheduler transmission rate for queues 0 through 7 are 95, 0, 0, 0, 0, 0, 0, and 5 percent of the total available bandwidth. The default buffer-size percentages for queues 0 through 7 are 95, 0, 0, 0, 0, 0, 0, and 5 percent of the total available buffer.

NOTE: On J-EX8200 switches, the default scheduler transmission rates for queues 0 through 7 are 75, 0, 20, 0, 0, 0, 0, and 5 percent of the total available bandwidth. The default buffer-size percentages for queues 0 through 7 are 75, 0, 20, 0, 0, 0, 0, and 5 percent of the total available buffer.

For each scheduler, you can configure the buffer size as one of the following:

- A percentage of the total buffer.
- The remaining buffer available. The remainder is the buffer percentage that is not assigned to other queues. For example, if you assign 40 percent of the delay buffer to queue 0, allow queue 2 to keep the default allotment of 20 percent, allow queue 7 to keep the default allotment of 5 percent, and assign the remainder to queue 3, then queue 3 uses approximately 35 percent of the delay buffer.

Priority Scheduling

Priority scheduling determines the order in which an output interface transmits traffic from the queues, thus ensuring that queues containing important traffic are provided better access to the outgoing interface.

Priority scheduling is accomplished through a procedure in which the scheduler examines the priority of the queue. The Junos OS supports two levels of transmission priority:

- Low—The scheduler determines if the individual queue is within its defined bandwidth profile. This binary decision, which is reevaluated on a regular time cycle, compares the amount of data transmitted by the queue against the amount of bandwidth allocated to it by the scheduler. When the transmitted amount is less than the allocated amount, the queue is considered to be in profile. A queue is out of profile when its transmitted amount is larger than its allocated amount. Out of profile queue will be transmitted only if bandwidth is available. Otherwise, it will be buffered.

  A queue from the set is selected based on the shaped deficit weighted round robin (SDWRR) algorithm, which operates within the set.

- Strict-high—Strict-high priority queue receives preferential treatment over low priority queue. Unlimited bandwidth is assigned to strict-high priority queue. Queues are scheduled according to the queue number, starting with the highest queue 7, with decreasing priority down through queue 0. Traffic in higher queue numbers is always scheduled prior to traffic in lower queue numbers. In other words, in case of two high priority queues, the queue with higher queue number is processed first.

Packets in low priority queues are transmitted only when strict-high priority queues are empty.

Scheduler Drop-Profile Maps

Drop-profile maps associate drop profiles with a scheduler. Drop-profile map sets the drop profile for a specific packet loss priority (PLP) and protocol type. The inputs for the drop-profile map are the PLP and the protocol type. The output is the drop profile.

Scheduler Maps

A scheduler map associates a specified forwarding class with a scheduler configuration. After configuring a scheduler, you must include it in a scheduler map and then associate the scheduler map with an output interface.

J-EX Series Switches allow you to associate up to four user-defined scheduler maps with interfaces.

Related Documentation

- Understanding Junos OS CoS Components for J-EX Series Switches on page 2862
- Example: Configuring CoS on J-EX Series Switches on page 2883
- Defining CoS Schedulers (CLI Procedure) on page 2920
- Defining CoS Schedulers (J-Web Procedure) on page 2920

Understanding CoS Two-Color Marking

Networks police traffic by limiting the input or output transmission rate of a class of traffic on the basis of user-defined criteria. Policing traffic allows you to control the maximum rate of traffic sent or received on an interface and to partition a network into multiple priority levels or classes of service.

Policers require you to apply limits to the traffic flow and set a consequence for packets that exceed these limits—usually a higher loss priority, so that packets exceeding the policer limits are discarded first.

J-EX Series Switches support a single-rate two-color marking type of policer, which is a simplified version of Single-Rate-Three-Color marking, defined in RFC 2697, *A Single Rate Three Color Marker*. This type of policer meters traffic based on the configured committed information rate (CIR) and committed burst size (CBS).

The single-rate two-color marker meters traffic and marks incoming packets depending on whether they are smaller than the committed burst size (CBS)—marked green—or exceed it— marked red.

The single-rate two-color marking policer operates in color-blind mode. In this mode, the policer's actions are not affected by any previous marking or metering of the examined packets. In other words, the policer is “blind? to any previous coloring a packet might have had.

Related Documentation

- Understanding Junos OS CoS Components for J-EX Series Switches on page 2862
- Understanding the Use of Policers in Firewall Filters on page 2752
- Configuring Policers to Control Traffic Rates (CLI Procedure) on page 2788

Understanding CoS Rewrite Rules

As packets enter or exit a network, edge switches might be required to alter the class-of-service (CoS) settings of the packets. This topic describes how to use rewrite rules to alter the CoS settings. It covers:

How Rewrite Rules Work

Rewrite rules set the value of the CoS bits within the packet’s header. Each rewrite rule reads the current forwarding class and loss priority associated with the packet, locates the chosen CoS value from a table, and writes this CoS value into the packet header. For rewrites to occur, rewrite rules must be explicitly assigned to an interface. Only tagged Layer 3 interfaces and tagged routed VLAN interfaces (RVIs) automatically rewrite packets by using the default IEEE 802.1p rewrite rule. Multiple rewrite rules of different types can be assigned to a single interface.

NOTE: On J-EX8200 Ethernet Switches, tagged Layer 3 interfaces and tagged RVIs do not automatically rewrite packets using the default IEEE 802.1p rewrite rule. You must explicitly assign the IEEE 802.1p rewrite rule to these interfaces for rewrites to occur.

Also, only one rewrite rule of each type can be assigned to any interface on a J-EX8200 switch.

In effect, the rewrite rule performs the opposite function of the behavior aggregate (BA) classifier used when the packet enters the switch. As the packet leaves the switch, the final CoS action is generally the application of a rewrite rule.

You configure rewrite rules to alter CoS values in outgoing packets on the outbound interfaces of an edge switch to meet the policies of a targeted peer. This allows the downstream switch in a neighboring network to classify each packet into the appropriate service group.

NOTE: When an IP precedence rewrite rule is active, bits 3, 4, and 5 of the ToS byte are always reset to zero when code points are rewritten.

Default Rewrite Rule

To enable a rewrite rule on an interface, you can either create your own rewrite rule and enable it on the interface or enable a default rewrite rule. See “Defining CoS Rewrite Rules (CLI Procedure)” on page 2925.

Table 363 on page 2877 shows the default rewrite-rule mappings. These are based on the default bit definitions of Differentiated Services code point (DSCP), IEEE 802.1p, and IP precedence values and the default forwarding classes.

When the CoS values of a packet match the forwarding-class and packet-loss-priority (PLP) values, the switch rewrites markings on the packet based on the rewrite table.

Table 363: Default Packet Header Rewrite Mappings

Map from Forwarding Class	PLP Value	Map to DSCP/IEEE 802.1p/IP Precedence Value
expedited-forwarding	low	ef
expedited-forwarding	high	ef
assured-forwarding	low	af11
assured-forwarding	high	af12 (DSCP)
best-effort	low	be

Table 363: Default Packet Header Rewrite Mappings *(continued)*

Map from Forwarding Class	PLP Value	Map to DSCP/IEEE 802.1p/IP Precedence Value
best-effort	high	be
network-control	low	nc1/cs6
network-control	high	nc2/cs7

Related Documentation

- Understanding Junos OS CoS Components for J-EX Series Switches on page 2862
- Example: Configuring CoS on J-EX Series Switches on page 2883
- Defining CoS Rewrite Rules (CLI Procedure) on page 2925
- Defining CoS Rewrite Rules (J-Web Procedure) on page 2926

Understanding Port Shaping and Queue Shaping for CoS on J-EX Series Switches

If the amount of traffic on a switch's network interface is more than the maximum bandwidth allowed on the interface, it leads to congestion. Port shaping and queue shaping can be used to manage the excess traffic and avoid congestion. Port shaping defines the maximum bandwidth allocated to a port, while queue shaping defines a limit on excess-bandwidth usage per queue.

This topic covers:

Port Shaping

Port shaping enables you to shape the aggregate traffic through a port or channel to a rate that is less than the line or port rate.

Queue Shaping

Queue shaping throttles the rate at which queues transmit packets. For example, using queue shaping, you can rate-limit a strict-priority queue so that the strict-priority queue does not lock out (or starve) low-priority queues. Similarly, for any queue, you can configure queue shaping.

Related Documentation

- Understanding CoS Schedulers on page 2873
- Defining CoS Schedulers (CLI Procedure) on page 2920

Understanding Junos OS EZQoS for CoS Configurations on J-EX Series Switches

Junos OS EZQoS on J-EX Series Switches eliminates the complexities involved in configuring class of service (CoS) across the network. EZQoS offers templates for key traffic classes.

Junos OS CoS allows you to divide traffic into classes and offer various levels of throughput and packet loss when congestion occurs. You can use CoS to ensure that different types of traffic (voice, video, and data) get the bandwidth and consideration they need to meet user expectations and business objectives.

Configuring CoS requires careful consideration of your service needs and thorough planning and design to ensure consistency across all switches in a CoS domain. To configure CoS manually, you must define and fine-tune all CoS components such as classifiers, rewrite rules, forwarding classes, schedulers, and scheduler-maps and then apply these components to the interfaces. Therefore, configuring CoS can be a fairly complex and time-consuming task.

EZQoS works by automatically assigning preconfigured values to all CoS parameters based on the typical application requirements. These preconfigured values are stored in a template with a unique name. You can change the preconfigured values of these parameters to suit your particular application needs.

For using EZQoS, you must identify which switch ports are being used for a specific application (such as VoIP, video, and data) and manually apply the corresponding application-specific EZQoS template to these switch ports.

NOTE: Currently, we provide an EZQoS template for configuring CoS for VoIP.

NOTE: We recommend that you do not use the term **EZQoS** for defining a classifier.

Related Documentation

- Junos OS CoS for J-EX Series Switches Overview on page 2860
- Configuring Junos OS EZQoS for CoS (CLI Procedure) on page 2930

Understanding Using CoS with MPLS Networks on J-EX Series Switches

You can use class of service (CoS) within MPLS networks to prioritize certain types of traffic during periods of congestion.

J-EX Series Switches support Differentiated Service Code Point (DSCP) or IP precedence and IEEE 802.1p CoS classifiers on the customer-edge interfaces of the ingress provider edge (PE) switch. DSCP or IP precedence classifiers are used for Layer 3 packets. IEEE 802.1p is used for Layer 2 packets.

When a packet enters a customer-edge interface of the ingress PE switch, the switch associates the packet with a particular CoS servicing level prior to putting the packet onto the label-switched path (LSP). The switches within the LSP utilize the CoS value set at the ingress PE switch. The CoS value that was embedded in the DSCP, IP precedence, or IEEE 802.1 p classifier is translated and encoded in the MPLS header by means of the EXP or experimental bits.

J-EX Series switches enable a default EXP classifier and a default EXP rewrite rule. You can configure a custom EXP classifier and a custom EXP rewrite rule if you prefer. However, the switch supports only one type of EXP classifier (default or custom) and only one EXP rewrite rule (default or custom).

You do not bind the EXP classifier or the EXP rewrite rule to individual interfaces. The switch automatically and implicitly applies the default or the custom EXP classifier and the default or the custom EXP rewrite rule to the appropriate MPLS-enabled interfaces. Because rewrite rules affect only egress interfaces, the switch applies the EXP rewrite rule only to those MPLS interfaces that are transmitting MPLS packets (not to the MPLS interfaces that are receiving the packets).

This topic includes:

Guidelines for Using CoS Classifiers on CCCs

When you are configuring CoS for MPLS over circuit cross-connect (CCC), there are some additional guidelines, as follows:

- You *must* explicitly bind a CoS classifier to the CCC interface on the ingress PE switch.
- You *cannot* use more than one type of DSCP/IP precedence and not more than one type of IEEE 802.1p classifier on the CCC interfaces. Thus, if you configure one CCC interface to use DSCP1, you cannot configure another CCC interface to use DSCP2. Likewise, if you configure one CCC interface to use IEEE1, you cannot configure another

CCC interface on the same switch to use IEEE2. All the CCC interfaces on the switch must use the same DSCP classifier and the same type of IEEE 802.1p classifier.

- You *cannot* configure one CCC interface as DSCP and another CCC interface as IP precedence, because these classifier types overlap.
- You *can* configure one CCC interface as DSCP and another CCC interface as IEEE 802.1p.
- You *can* configure one CCC interface as both DSCP and IEEE 802.1p. If you configure a CCC interface with both these classifiers, the DSCP classifier is used for routing Layer 3 packets and the IEEE 802.1p classifier is used for routing Layer 2 packets.

NOTE: You can define multiple types of DSCP, IP precedence, and IEEE 802.1p on the switch and use the different classifier types for the non-CCC interfaces on the switch.

Using CoS Classifiers with IP over MPLS

When you are configuring CoS for IP over MPLS, the customer-edge interface uses the CoS configuration that has been set up for the switch as the default. You do not have to bind a classifier to the customer-edge interface in this case. There are no restrictions regarding using multiple types of DSCP, IP precedence, and IEEE 802.1p on the same switch.

- You can modify the CoS classifier for a particular interface, but it is not required.
- You can configure one interface as DSCP1 and another as DSCP2 and another and IP precedence, and so forth.

Default Classifiers and Default Rewrite Rules

The default classifiers support only two forwarding classes, **best-effort** and **network-control**, and use only two queues, **0** and **7**. However, J-EX Series switches support up to sixteen forwarding classes and eight queues. To use the additional forwarding classes and queues, create a custom classifier. To modify the code point and loss priority for a specific forwarding class, configure a rewrite rule on the switch. The default rewrite rule for EXP is enabled in the default configuration. However, the default rewrite rules for the other classifiers are not enabled in the default configuration. You can display the default classifier mappings and default rewrite mappings by entering the **show class-of-service** command on the switch.

EXP Rewrite Rules

When traffic passes from the customer-edge interface to an MPLS interface, the DSCP, IP precedence, or IEEE 802.1p CoS classifier is translated into the EXP bits within the MPLS header. You cannot disable the default EXP rewrite rule, but you can configure your own custom EXP classifier and a custom EXP rewrite rule. You cannot bind the EXP classifier to individual MPLS interfaces; the switch applies it globally to all the MPLS-enabled interfaces on the switch.

Only one EXP rewrite rule (either default or custom) is supported on a switch. The switch applies it to all the MPLS-enabled egress interfaces.

Policer

Policing helps to ensure that the amount of traffic forwarded through an LSP never exceeds the requested bandwidth allocation. During periods of congestion (when the total rate of queuing packets exceeds the rate of transmission), any new packets being sent to an interface can be dropped because there is no place to store them. You should configure a policer on the ingress PE switch:

- If you are using MPLS with CCC, you bind the policer to the LSP. You cannot bind a policer to a CCC interface.
- If you are using IP over MPLS, you bind the policer to the **inet-family** customer-edge interface. You cannot bind a policer to the LSP when you are using IP over MPLS.

Schedulers

The schedulers for using CoS with MPLS are the same as for the other CoS configurations on J-EX Series switches. Default schedulers are provided for **best-effort** and **network-control** forwarding classes. If you are using **assured-forwarding**, **expedited-forwarding**, or other custom forwarding classes, we recommend that you configure a scheduler to support that forwarding class. See “Understanding CoS Schedulers” on page 2873.

Related Documentation

- Junos OS MPLS for J-EX Series Switches Overview on page 3057
- Understanding CoS Classifiers on page 2867
- Understanding CoS Schedulers on page 2873
- Example: Configuring CoS on J-EX Series Switches on page 2883
- Configuring CoS on MPLS Provider Edge Switch Using Circuit Cross-Connect (CLI Procedure) on page 2932
- Configuring Rewrite Rules for EXP Classifiers on MPLS Networks (CLI Procedure)
- Configuring CoS on Provider Switches of an MPLS Network (CLI Procedure) on page 3106
- Defining CoS Rewrite Rules (CLI Procedure) on page 2925
- Configuring Policers to Control Traffic Rates (CLI Procedure) on page 2788

CHAPTER 108

Examples: CoS Configuration

Example: Configuring CoS on J-EX Series Switches

Configure class of service (CoS) on your switch to manage traffic so that when the network experiences congestion and delay, critical applications are protected. Using CoS, you can divide traffic on your switch into classes and provide various levels of throughput and packet loss. This is especially important for traffic that is sensitive to jitter and delay, such as voice traffic.

This example shows how to configure CoS on a single J-EX Series switch in the network.

Requirements

This example uses the following hardware and software components:

- One J-EX4200 switch

Overview and Topology

This example uses the topology shown in Figure 80 on page 2884.

Figure 80: Topology for Configuring CoS

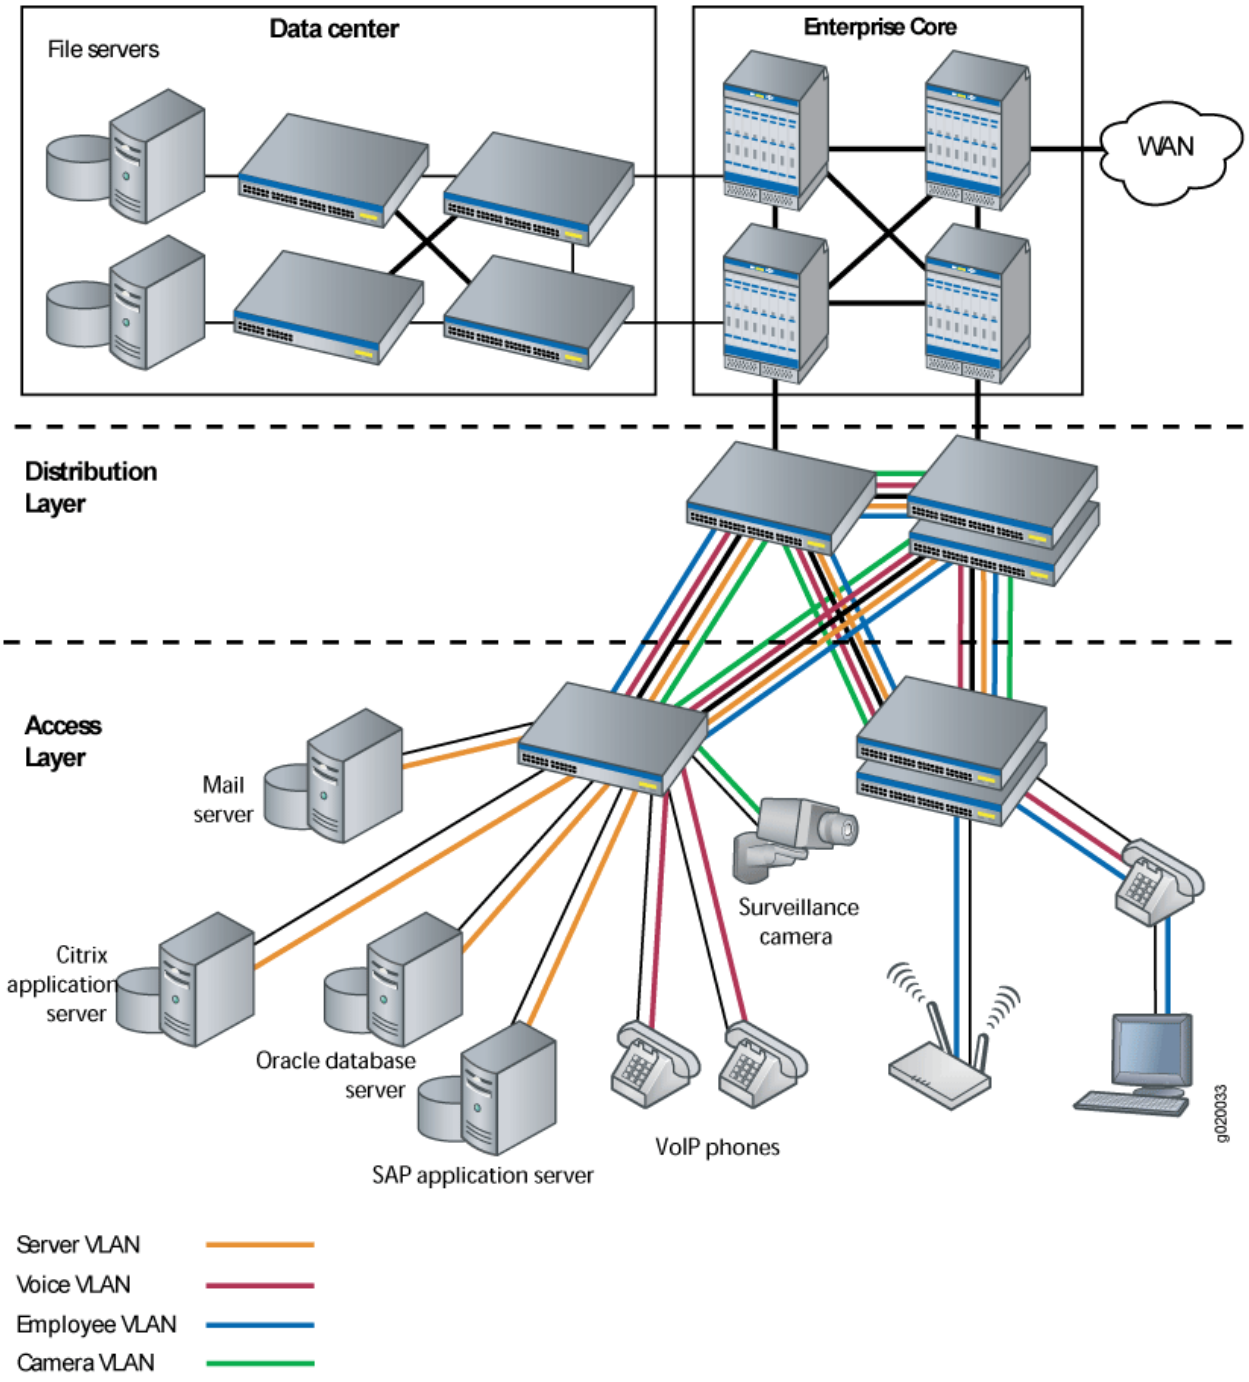

The topology for this configuration example consists of one J-EX Series switch at the access layer.

The J-EX Series access switch is configured to support VLAN membership. Switch ports **ge-0/0/0**and **ge-0/0/1** are assigned to the **voice-vlan** for two VoIP phones. Switch port **ge-0/0/2** is assigned to the **camera-vlan** for the surveillance camera. Switch ports

ge-0/0/3, **ge-0/0/4**, **ge-0/0/5**, and **ge-0/0/6** are assigned to the **server-vlan** for the servers hosting various applications such as those provided by Citrix, Microsoft, Oracle, and SAP.

Table 364 on page 2885 shows the VLAN configuration components.

Table 364: Configuration Components: VLANs

VLAN Name	VLAN ID	VLAN Subnet and Available IP Addresses	VLAN Description
voice-vlan	10	**192.168.1.0/32** **192.168.1.1** through **192.168.1.11** **192.168.1.12** is the subnet's broadcast address.	Voice VLAN used for employee VoIP communication.
camera-vlan	20	**192.168.1.13/32** **192.168.1.14** through **192.168.1.20** **192.168.1.21** is the subnet's broadcast address.	VLAN for the surveillance cameras.
server-vlan	30	**192.168.1.22/32** **192.168.1.23** through **192.168.1.35** **192.168.1.36** is the subnet's broadcast address.	VLAN for the servers hosting enterprise applications.

Ports on the J-EX Series switches support Power over Ethernet (PoE) to provide both network connectivity and power for VoIP telephones connecting to the ports. Table 365 on page 2885 shows the switch interfaces that are assigned to the VLANs and the IP addresses for devices connected to the switch ports:

Table 365: Configuration Components: Switch Ports on a 48-Port All-PoE Switch

Interfaces	VLAN Membership	IP Addresses	Port Devices
ge-0/0/0, **ge-0/0/1**	**voice-vlan**	**192.168.1.1** through **192.168.1.2**	Two VoIP telephones.
ge-0/0/2	**camera-vlan**	**192.168.1.14**	Surveillance camera.
ge-0/0/3, **ge-0/0/4**, **ge-0/0/5**, **ge-0/0/6**	**sevrer-vlan**	**192.168.1.23** through **192.168.1.26**	Four servers hosting applications such as those provided by Citrix, Microsoft, Oracle, and SAP.

NOTE: This example shows how to configure CoS on a single J-EX Series switch. This example does not consider across-the-network applications of CoS in which you might implement different configurations on ingress and egress switches to provide differentiated treatment to different classes across a set of nodes in a network.

Configuration

CLI Quick Configuration

To quickly configure CoS, copy the following commands and paste them into the switch terminal window:

```
[edit]
set class-of-service forwarding-classes class app queue-num 5
set class-of-service forwarding-classes class mail queue-num 1
set class-of-service forwarding-classes class db queue-num 2
set class-of-service forwarding-classes class erp queue-num 3
set class-of-service forwarding-classes class video queue-num 4
set class-of-service forwarding-classes class best-effort queue-num 0
set class-of-service forwarding-classes class voice queue-num 6
set class-of-service forwarding-classes class network-control queue-num 7
set firewall family ethernet-switching filter voip_class term voip from source-address 192.168.1.1/32
set firewall family ethernet-switching filter voip_class term voip from source-address 192.168.1.2/32
set firewall family ethernet-switching filter voip_class term voip from protocol udp
set firewall family ethernet-switching filter voip_class term voip from source-port 2698
set firewall family ethernet-switching filter voip_class term voip then forwarding-class voice
loss-priority low
set firewall family ethernet-switching filter voip_class term network_control from precedence
[net-control internet-control]
set firewall family ethernet-switching filter voip_class term network_control then forwarding-class
network-control loss-priority low
set firewall family ethernet-switching filter voip_class term best_effort_traffic then
forwarding-class best-effort loss-priority low
set interfaces ge-0/0/0 description phone1—voip-ingress-port
set interfaces ge-0/0/0 unit 0 family ethernet-switching filter input voip_class
set interfaces ge-0/0/1 description phone2—voip-ingress-port
set interfaces ge-0/0/1 unit 0 family ethernet-switching filter input voip_class
set firewall family ethernet-switching filter video_class term video from source-address
192.168.1.14/32
set firewall family ethernet-switching filter video_class term video from protocol udp
set firewall family ethernet-switching filter video_class term video from source-port 2979
set firewall family ethernet-switching filter video_class term video then forwarding-class video
loss-priority low
set firewall family ethernet-switching filter video_class term network_control from precedence
[net-control internet-control]
set firewall family ethernet-switching filter video_class term network_control then forwarding-class
network-control loss-priority low
set firewall family ethernet-switching filter video_class term best_effort_traffic then
forwarding-class best-effort loss-priority low
set interfaces ge-0/0/2 description video-ingress-port
set interfaces ge-0/0/2 unit 0 family ethernet-switching filter input video_class
set firewall family ethernet-switching filter app_class term app from source-address
192.168.1.23/32
set firewall family ethernet-switching filter app_class term app from protocol tcp
set firewall family ethernet-switching filter app_class term app from source-port [1494 2512 2513
2598 2897]
```

```
set firewall family ethernet-switching filter app_class term app then forwarding-class app
loss-priority low
set firewall family ethernet-switching filter app_class term mail from source-address
192.168.1.24/32
set firewall family ethernet-switching filter app_class term mail from protocol tcp
set firewall family ethernet-switching filter app_class term mail from source-port [25 143 389
691 993 3268 3269]
set firewall family ethernet-switching filter app_class term mail then forwarding-class mail
loss-priority low
set firewall family ethernet-switching filter app_class term db from source-address 192.168.1.25/32
set firewall family ethernet-switching filter app_class term db from protocol tcp
set firewall family ethernet-switching filter app_class term db from source-port [1521 1525 1527
1571 1810 2481]
set firewall family ethernet-switching filter app_class term db then forwarding-class db loss-priority
low
set firewall family ethernet-switching filter app_class term erp from source-address 192.168.1.26/32
set firewall family ethernet-switching filter app_class term erp from protocol tcp
set firewall family ethernet-switching filter app_class term erp from source-port [3200 3300
3301 3600]
set firewall family ethernet-switching filter app_class term erp then forwarding-class erp
loss-priority low
set firewall family ethernet-switching filter app_class term network_control from precedence
[net-control internet-control]
set firewall family ethernet-switching filter app_class term network_control then forwarding-class
network-control loss-priority low
set firewall family ethernet-switching filter app_class term best_effort_traffic then forwarding-class
best-effort loss-priority low
set interfaces ge-0/0/3 unit 0 family ethernet-switching filter input app_class
set interfaces ge-0/0/4 unit 0 family ethernet-switching filter input app_class
set interfaces ge-0/0/5 unit 0 family ethernet-switching filter input app_class
set interfaces ge-0/0/6 unit 0 family ethernet-switching filter input app_class
set class-of-service schedulers voice-sched buffer-size percent 10
set class-of-service schedulers voice-sched priority strict-high
set class-of-service schedulers voice-sched transmit-rate percent 10
set class-of-service schedulers video-sched buffer-size percent 15
set class-of-service schedulers video-sched priority low
set class-of-service schedulers video-sched transmit-rate percent 15
set class-of-service schedulers app-sched buffer-size percent 10
set class-of-service schedulers app-sched priority low
set class-of-service schedulers app-sched transmit-rate percent 10
set class-of-service schedulers mail-sched buffer-size percent 5
set class-of-service schedulers mail-sched priority low
set class-of-service schedulers mail-sched transmit-rate percent 5
set class-of-service schedulers db-sched buffer-size percent 10
set class-of-service schedulers db-sched priority low
set class-of-service schedulers db-sched transmit-rate percent 10
set class-of-service schedulers erp-sched buffer-size percent 10
set class-of-service schedulers erp-sched priority low
set class-of-service schedulers erp-sched transmit-rate percent 10
set class-of-service schedulers nc-sched buffer-size percent 5
set class-of-service schedulers nc-sched priority strict-high
set class-of-service schedulers nc-sched transmit-rate percent 5
set class-of-service schedulers be-sched buffer-size percent 35
set class-of-service schedulers be-sched priority low
set class-of-service schedulers be-sched transmit-rate percent 35
set class-of-service scheduler-maps ethernet-cos-map forwarding-class voice scheduler
voice-sched
set class-of-service scheduler-maps ethernet-cos-map forwarding-class video scheduler
video-sched
set class-of-service scheduler-maps ethernet-cos-map forwarding-class app scheduler app-sched
```

```
set class-of-service scheduler-maps ethernet-cos-map forwarding-class mail scheduler mail-sched
set class-of-service scheduler-maps ethernet-cos-map forwarding-class db scheduler db-sched
set class-of-service scheduler-maps ethernet-cos-map forwarding-class erp scheduler erp-sched
set class-of-service scheduler-maps ethernet-cos-map forwarding-class network-control
scheduler nc-sched
set class-of-service scheduler-maps ethernet-cos-map forwarding-class best-effort scheduler
be-sched
set class-of-service interfaces ge-0/0/20 scheduler-map ethernet-cos-map
```

Step-by-Step Procedure

To configure and apply CoS:

1. Configure one-to-one mapping between eight forwarding classes and eight queues:

```
[edit class-of-service]
user@switch# set forwarding-classes class app queue-num 5
user@switch# set forwarding-classes class mail queue-num 1
user@switch# set forwarding-classes class db queue-num 2
user@switch# set forwarding-classes class erp queue-num 3
user@switch# set forwarding-classes class video queue-num 4
user@switch# set forwarding-classes class best-effort queue-num 0
user@switch# set forwarding-classes class voice queue-num 6
user@switch# set forwarding-classes class network-control queue-num 7
```

2. Define the firewall filter **voip_class** to classify the VoIP traffic:

```
[edit firewall]
user@switch# set family ethernet-switching filter voip_class
```

3. Define the term **voip**:

```
[edit firewall]
user@switch# set family ethernet-switching filter voip_class term voip from
source-address 192.168.1.1/32
user@switch# set family ethernet-switching filter voip_class term voip from
source-address 192.168.1.2/32
user@switch# set family ethernet-switching filter voip_class term voip protocol udp
user@switch# set family ethernet-switching filter voip_class term voip source-port
2698
user@switch# set family ethernet-switching filter voip_class term voip then
forwarding-class voice loss-priority low
```

4. Define the term **network_control**:

```
[edit firewall]
user@switch# set family ethernet-switching filter voip_class term network_control from
precedence [net-control internet-control]
user@switch# set family ethernet-switching filter voip_class term network_control then
forwarding-class network-control loss-priority low
```

5. Define the term **best_effort_traffic** with no match conditions:

```
[edit firewall]
user@switch# set family ethernet-switching filter voip_class term best_effort_traffic
then forwarding-class best-effort loss-priority low
```

6. Apply the firewall filter **voip_class** as an input filter to the interfaces for the VoIP phones:

```
[edit interfaces]
user@switch# set ge-0/0/0 description phone1—voip-ingress-port
user@switch# set ge-0/0/0 unit 0 family ethernet-switching filter input voip_class
user@switch# set ge-0/0/1 description phone2—voip-ingress-port
```

```
user@switch# set ge-0/0/1 unit 0 family ethernet-switching filter input voip_class
```

7. Define the firewall filter **video_class** to classify the video traffic:

```
[edit firewall]
user@switch# set family ethernet-switching filter video_class
```

8. Define the term **video**:

```
[edit firewall]
user@switch# set family ethernet-switching filter video_class term video from
source-address 192.168.1.14/32
user@switch# set family ethernet-switching filter video_class term video protocol udp
user@switch# set family ethernet-switching filter video_class term video source-port
2979
user@switch# set family ethernet-switching filter video_class term video then
forwarding-class video loss-priority low
```

9. Define the term **network_control** (for the **video_class** filter):

```
[edit firewall]
user@switch# set family ethernet-switching filter video_class term network_control
from precedence [net-control internet-control]
user@switch# set family ethernet-switching filter video_class term network_control
then forwarding-class network-control loss-priority low
```

10. Define the term **best_effort_traffic** (for the **video_class** filter):

```
[edit firewall]
user@switch# set family ethernet-switching filter video_class term best_effort_traffic
then forwarding-class best-effort loss-priority low
```

11. Apply the firewall filter **video_class** as an input filter to the interface for the surveillance camera:

```
[edit interfaces]
user@switch# set ge-0/0/2 description video-ingress-port
user@switch# set ge-0/0/2 unit 0 family ethernet-switching filter input video_class
```

12. Define the firewall filter **app_class** to classify the application server traffic:

```
[edit firewall]
user@switch# set family ethernet-switching filter app_class
```

13. Define the term **app**:

```
[edit firewall]
user@switch# set family ethernet-switching filter app_class term app from
source-address 192.168.1.23/32
user@switch# set family ethernet-switching filter app_class term app protocol tcp
user@switch# set family ethernet-switching filter app_class term app source-port [1494
2512 2513 2598 2897]
user@switch# set family ethernet-switching filter app_class term app then
forwarding-class app loss-priority low
```

14. Define the term **mail**:

```
[edit firewall]
user@switch# set family ethernet-switching filter app_class term mail from
source-address 192.168.1.24/32
user@switch# set family ethernet-switching filter app_class term mail protocol tcp
user@switch# set family ethernet-switching filter app_class term mail source-port [25
143 389 691 993 3268 3269]
```

```
user@switch# set family ethernet-switching filter app_class term mail then
forwarding-class mail loss-priority low
```

15. Define the term **db**:

```
[edit firewall]
user@switch# set family ethernet-switching filter app_class term db from source-address
192.168.1.25/32
user@switch# set family ethernet-switching filter app_class term db protocol tcp
user@switch# set family ethernet-switching filter app_class term db source-port [1521
1525 1527 1571 1810 2481]
user@switch# set family ethernet-switching filter app_class term db then
forwarding-class db loss-priority low
```

16. Define the term **erp**:

```
[edit firewall]
user@switch# set family ethernet-switching filter app_class term erp from
source-address 192.168.1.26/32
user@switch# set family ethernet-switching filter app_class term erp protocol tcp
user@switch# set family ethernet-switching filter app_class term erp source-port [3200
3300 3301 3600]
user@switch# set family ethernet-switching filter app_class term erp then
forwarding-class erp loss-priority low
```

17. Define the term **network_control** (for the **app_class** filter):

```
[edit firewall]
user@switch# set family ethernet-switching filter app_class term network_control from
precedence [net-control internet-control]
user@switch# set family ethernet-switching filter app_class term network_control then
forwarding-class network-control loss-priority low
```

18. Define the term **best_effort_traffic** (for the **app_class** filter):

```
[edit firewall]
user@switch# set family ethernet-switching filter app_class term best_effort_traffic
then forwarding-class best-effort loss-priority low
```

19. Apply the firewall filter **app_class** as an input filter to the interfaces for the servers hosting applications:

```
[edit interfaces]
user@switch# set ge-0/0/3 unit 0 family ethernet-switching filter input app_class
user@switch# set ge-0/0/4 unit 0 family ethernet-switching filter input app_class
user@switch# set ge-0/0/5 unit 0 family ethernet-switching filter input app_class
user@switch# set ge-0/0/6 unit 0 family ethernet-switching filter input app_class
```

20. Configure schedulers:

```
[edit class-of-service]
user@switch# set schedulers voice-sched buffer-size percent 10
user@switch# set schedulers voice-sched priority strict-high
user@switch# set schedulers voice-sched transmit-rate percent 10
user@switch# set schedulers video-sched buffer-size percent 15
user@switch# set schedulers video-sched priority low
user@switch# set schedulers video-sched transmit-rate percent 15
user@switch# set schedulers app-sched buffer-size percent 10
user@switch# set schedulers app-sched priority low
user@switch# set schedulers app-sched transmit-rate percent 10
user@switch# set schedulers mail-sched buffer-size percent 5
user@switch# set schedulers mail-sched priority low
```

```
user@switch# set schedulers mail-sched transmit-rate percent 5
user@switch# set schedulers db-sched buffer-size percent 10
user@switch# set schedulers db-sched priority low
user@switch# set schedulers db-sched transmit-rate percent 10
user@switch# set schedulers erp-sched buffer-size percent 10
user@switch# set schedulers erp-sched priority low
user@switch# set schedulers erp-sched transmit-rate percent 10
user@switch# set schedulers nc-sched buffer-size percent 5
user@switch# set schedulers nc-sched priority strict-high
user@switch# set schedulers nc-sched transmit-rate percent 5
user@switch# set schedulers be-sched buffer-size percent 35
user@switch# set schedulers be-sched priority low
user@switch# set schedulers be-sched transmit-rate percent 35
```

21. Assign the forwarding classes to schedulers with the scheduler map **ethernet-cos-map**:

```
[edit class-of-service]
user@switch# set scheduler-maps ethernet-cos-map forwarding-class voice scheduler
voice-sched
user@switch# set scheduler-maps ethernet-cos-map forwarding-class video scheduler
video-sched
user@switch# set scheduler-maps ethernet-cos-map forwarding-class app scheduler
app-sched
user@switch# set scheduler-maps ethernet-cos-map forwarding-class mail scheduler
mail-sched
user@switch# set scheduler-maps ethernet-cos-map forwarding-class db scheduler
db-sched
user@switch# set scheduler-maps ethernet-cos-map forwarding-class erp scheduler
erp-sched
user@switch# set scheduler-maps ethernet-cos-map forwarding-class network-control
scheduler nc-sched
user@switch# set scheduler-maps ethernet-cos-map forwarding-class best-effort
scheduler be-sched
```

22. Associate the scheduler map with the outgoing interface:

```
[edit class-of-service interfaces]
user@switch# set ge-0/0/20 scheduler-map ethernet-cos-map
```

Results Display the results of the configuration:

```
user@switch# show firewall

firewall family ethernet-switching {
    filter voip_class {
        term voip {
            from {
                source-address {
                    192.168.1.1/32;
                    192.168.1.2/32;
                }
                protocol udp;
                source-port 2698;
            }
            then {
                forwarding-class voice;
                loss-priority low;
            }
```

```
    }
    term network control {
      from {
        precedence [net-control internet-control];
      }
      then {
        forwarding-class network-control;
        loss-priority low;
      }
    }
    term best_effort_traffic {
      then {
        forwarding-class best-effort;
        loss-priority low;
      }
    }
}
filter video_class {
    term video {
      from {
        source-address {
          192.168.1.14/32;
        }
        protocol udp;
        source-port 2979;
      }
      then {
        forwarding-class video;
        loss-priority low;
      }
    }
    term network control {
      from {
        precedence [net-control internet-control];
      }
      then {
        forwarding-class network-control;
        loss-priority low;
      }
    }
    term best_effort_traffic {
      then {
        forwarding-class best-effort;
        loss-priority low;
      }
    }
}
filter app_class {
    term app {
      from {
        source-address {
          192.168.1.23/32;
        }
        protocol tcp;
        source-port [1491 2512 2513 2598 2897];
      }
```

```
    then {
      forwarding-class app;
      loss-priority low;
    }
  }
  term mail {
    from {
      source-address {
        192.168.1.24/32;
      }
      protocol tcp;
      source-port [25 143 389 691 993 3268 3269];
    }
    then {
      forwarding-class mail;
      loss-priority low;
    }
  }
  term db {
    from {
      source-address {
        192.168.1.25/32;
      }
      protocol tcp;
      source-port [1521 1525 1527 1571 1810 2481];
    }
    then {
      forwarding-class db;
      loss-priority low;
    }
  }
  term erp {
    from {
      source-address {
        192.168.1.26/32;
      }
      protocol tcp;
      source-port [3200 3300 3301 3600];
    }
    then {
      forwarding-class erp;
      loss-priority low;
    }
  }
  term network control {
    from {
      precedence [net-control internet-control];
    }
    then {
      forwarding-class network-control;
      loss-priority low;
    }
  }
  term best_effort_traffic {
    then {
      forwarding-class best-effort;
```

```
            loss-priority low;
          }
        }
      }
    }
}

user@switch# show class-of-service

forwarding-classes {
    class app queue-num 5;
    class mail queue-num 1;
    class db queue-num 2;
    class erp queue-num 3;
    class video queue-num 4;
    class best-effort queue-num 0;
    class voice queue-num 6;
    class network-control queue-num 7;
}
schedulers {
    voice-sched {
        buffer-size percent 10;
        priority strict-high;
        transmit-rate percent 10;
    }
    video-sched {
        buffer-size percent 15;
        priority low;
        transmit-rate percent 15;
    }
    app-sched {
        buffer-size percent 10;
        priority low;
        transmit-rate percent 10;
    }
    mail-sched {
        buffer-size percent 5;
        priority low;
        transmit-rate percent 5;
    }
    db-sched {
        buffer-size percent 10;
        priority low;
        transmit-rate percent 10;
    }
    erp-sched {
        buffer-size percent 10;
        priority low;
        transmit-rate percent 10;
    }
    nc-sched {
        buffer-size percent 5;
        priority strict-high;
        transmit-rate percent 5;
    }
    be-sched {
        buffer-size percent 35;
        priority low;
```

```
      transmit-rate percent 35;
    }
  }
  scheduler-maps {
    ethernet-cos-map {
      forwarding-class voice scheduler voice-sched;
      forwarding-class video scheduler video-sched;
      forwarding-class app scheduler app-sched;
      forwarding-class mail scheduler mail-sched;
      forwarding-class db scheduler db-sched;
      forwarding-class erp scheduler erp-sched;
      forwarding-class network-control scheduler nc-sched;
      forwarding-class best-effort scheduler be-sched;
    }
  }
```

```
user@switch# show interfaces
```

```
ge-0/0/0 {
    unit 0 {
        family ethernet {
            filter {
                input voip_class;
            }
        }
    }
}
ge-0/0/1 {
    unit 0 {
        family ethernet {
            filter {
                input voip_class;
            }
        }
    }
}
ge-0/0/2 {
    unit 0 {
        family ethernet {
            filter {
                input video_class;
            }
        }
    }
}
ge-0/0/3 {
    unit 0 {
        family ethernet {
            filter {
                input app_class;
            }
        }
    }
}
ge-0/0/4 {
    unit 0 {
        family ethernet {
```

```
        filter {
          input app_class;
        }
      }
    }
  }
  ge-0/0/5 {
    unit 0 {
      family ethernet {
        filter {
          input app_class;
        }
      }
    }
  }
  ge-0/0/6 {
    unit 0 {
      family ethernet {
        filter {
          input app_class;
        }
      }
    }
  }
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying That the Defined Forwarding Classes Exist and Are Mapped to Queues on page 2896
- Verifying That the Forwarding Classes Have Been Assigned to Schedulers on page 2897
- Verifying That the Scheduler Map Has Been Applied to the Interface on page 2898

Verifying That the Defined Forwarding Classes Exist and Are Mapped to Queues

Purpose Verify that the following forwarding classes **app**, **db**, **erp**, **mail**, **video**, and **voice** have been defined and mapped to queues.

Action

```
user@switch> show class-of-service forwarding-class
Forwarding class                          ID    Queue
  app                                      0        5
  db                                       1        2
  erp                                      2        3
  best-effort                              3        0
  mail                                     4        1
  voice                                    5        6
  video                                    6        4
  network-control                          7        7
```

Meaning This output shows that the forwarding classes have been defined and mapped to appropriate queues.

Verifying That the Forwarding Classes Have Been Assigned to Schedulers

Purpose Verify that the forwarding classes have been assigned to schedulers.

Action

```
user@switch> show class-of-service scheduler-map
Scheduler map: ethernet-cos-map, Index: 2
  Scheduler: voice-sched, Forwarding class: voice, Index: 22
    Transmit rate: 5 percent, Rate Limit: none, Buffer size: 15 percent,
    Priority: Strict-high
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

  Scheduler: video-sched, Forwarding class: video, Index: 22
    Transmit rate: 10 percent, Rate Limit: none, Buffer size: 10 percent,
    Priority: low
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

  Scheduler: app-sched, Forwarding class: app, Index: 22
    Transmit rate: 10 percent, Rate Limit: none, Buffer size: 10 percent,
    Priority: low
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

  Scheduler: mail-sched, Forwarding class: mail, Index: 22
    Transmit rate: 5 percent, Rate Limit: none, Buffer size: 5 percent,
    Priority: low
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

  Scheduler: db-sched, Forwarding class: db, Index: 22
    Transmit rate: 10 percent, Rate Limit: none, Buffer size: 10 percent,
    Priority: low
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

  Scheduler: erp-sched, Forwarding class: erp, Index: 22
    Transmit rate: 10 percent, Rate Limit: none, Buffer size: 10 percent,
    Priority: low
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

  Scheduler: be-sched, Forwarding class: best-effort, Index: 20
    Transmit rate: 35 percent, Rate Limit: none, Buffer size: 35 percent,
    Priority: low
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
```

```
    High            TCP             1     <default-drop-profile>

Scheduler: nc-sched, Forwarding class: network-control, Index: 22
  Transmit rate: 5 percent, Rate Limit: none, Buffer size: 5 percent,
  Priority: Strict-high
  Drop profiles:
    Loss priority   Protocol    Index    Name
    High            non-TCP         1    <default-drop-profile>
    High            TCP             1    <default-drop-profile>
```

Meaning This output shows that the forwarding classes have been assigned to schedulers.

Verifying That the Scheduler Map Has Been Applied to the Interface

Purpose Verify that the scheduler map has been applied to the interface.

Action

```
user@switch> show class-of-service interface
...
Physical interface: ge-0/0/20, Index: 149
Queues supported: 8, Queues in use: 8
  Scheduler map: ethernet-cos-map, Index: 43366
  Input scheduler map: <default>, Index: 3
...
```

Meaning This output shows that the scheduler map (**ethernet-cos-map**) has been applied to the interface (**ge-0/0/20**).

Related Documentation

- Defining CoS Code-Point Aliases (CLI Procedure) on page 2914
- Defining CoS Classifiers (CLI Procedure) on page 2914
- Defining CoS Forwarding Classes (CLI Procedure) on page 2918
- Defining CoS Schedulers (CLI Procedure) on page 2920
- Configuring CoS Tail Drop Profiles (CLI Procedure) on page 2925
- Assigning CoS Components to Interfaces (CLI Procedure) on page 2928
- Configuring Firewall Filters (CLI Procedure) on page 2779

Example: Combining CoS with MPLS on J-EX Series Switches

You can use class of service (CoS) within MPLS networks to prioritize certain types of traffic during periods of congestion. The CoS value is included within the MPLS label, which is passed through the network, enabling end-to-end CoS across the network.

MPLS services are often used to ensure better performance for low-latency applications such as VoIP and other business-critical functions. These applications place specific demands on a network for successful transmission. CoS gives you the ability to control the mix of bandwidth, delay, jitter, and packet loss while taking advantage of the MPLS labeling mechanism.

This example shows how to configure CoS on an MPLS network that is using a unidirectional circuit cross-connect (CCC) from the ingress provider edge (PE) switch to the egress PE switch. for the customer-edge interface of the ingress provider edge (PE) switch. It describes adding the configuration of CoS components to the ingress PE switch, the egress PE switch, and the core provider switches of the existing MPLS network. Because of the unidirectional configuration, the DSCP classifier needs to be configured only on the ingress PE switch.

Requirements

This example uses the following hardware and software components:

- Three J-EX Series switches

Before you configure CoS with MPLS, be sure you have:

Configured an MPLS network with two PE switches and one provider switch. See "Example: Configuring MPLS on J-EX Series Switches" on page 3071. This example assumes that an MPLS network has been configured using a cross circuit-connect (CCC).

Overview and Topology

This example describes adding custom classifiers and custom rewrite rules to switches in an MPLS network that is using MPLS over CCC.

It is a unidirectional configuration. Therefore, you need to configure custom classifiers and custom rewrite rules as follows:

- On the ingress PE switch: custom DSCP classifier and custom EXP rewrite rule
- On the egress PE switch: custom EXP classifier
- On the provider switch: customer EXP classifier and custom EXP rewrite rule

NOTE: You can also configure schedulers and shapers as needed. If you are using assured-forwarding, expedited-forwarding, or other custom forwarding classes, we recommend that you configure a scheduler to support that forwarding class. See "Defining CoS Schedulers (CLI Procedure)" on page 2920.

The example creates a custom DSCP classifier (**dscp1**) on the ingress PE switch and binds this classifier to the CCC interface. It includes configuration of a policer on the ingress PE switch. The policer is applied as a filter on the label-switched path (LSP) **lsp_to_pe2_ge1**(created in "Example: Configuring MPLS on J-EX Series Switches" on

page 3071) to ensure that the amount of traffic forwarded through the LSP never exceeds the requested bandwidth allocation.

This example creates a custom EXP rewrite rule (**exp1**) on the ingress PE switch, specifying a loss-priority and code point to be used for the expedited-forwarding class as the packet travels through the LSP. The switch applies this custom rewrite rule on the core interfaces **ge-0/0/5.0** and **ge-0/0/6.0**, which are the egress interfaces for this switch.

Table 366 on page 2900 shows the CoS configuration components added to the ingress PE switch.

Table 366: CoS Configuration Components on the Ingress PE Switch

Property	Settings	Description
Local PE switch hardware	J-EX Series switch	PE-1
Policing filter configured and applied to the LSP.	**policing filter mypolicer** **filter myfilter**	Name of the rate-limiting policer. Name of the filter, which refers to the policer
Custom DSCP classifier	**dscp1**	Specifies the name of the custom DSCP classifier
Custom EXP rewrite rule	**e1**	Name of the custom EXP rewrite rule.
Customer-edge interface	**ge-0/0/1.0**	Interface that receives packets from devices outside the network. The custom DSCP classifier must be specified on this CCC interface.
Core interfaces	**ge-0/0/5.0** and **ge-0/0/6.0**	Interfaces that transmit MPLS packets to other switches within the MPLS network. The EXP rewrite rule is applied implicitly to these interfaces.

Table 367 on page 2900 shows the CoS configuration components added to the egress PE switch in this example.

Table 367: CoS Configuration Components of the Egress PE Switch

Property	Settings	Description
Remote provider edge switch hardware	J-EX Series switch	PE-2
Custom EXP classifier	**exp1**	Name of custom EXP classifier

Table 367: CoS Configuration Components of the Egress PE Switch *(continued)*

Property	Settings	Description
Customer-edge interface	**ge-0/0/1.0**	Interface that transmits packets from this network to devices outside the network. No CoS classifier is specified for this interface. A scheduler can be specified.
Core interfaces	**ge-0/0/7.0** and **ge-0/0/8.0**	Core interfaces on PE-2 that receive MPLS packets from the provider switch. The EXP classifier is enabled by default on the switch and applied implicitly to these interfaces.

Table 368 on page 2901 shows the MPLS configuration components used for the provider switch in this example.

Table 368: CoS Configuration Components of the Provider Switch

Property	Settings	Description
Provider switch hardware	J-EX Series switch	Transit switch within the MPLS network configuration.
Custom EXP classifier	**exp1**	Name of the custom EXP classifier.
Custom EXP rewrite rule	**e1**	Name of the custom EXP rewrite rule.
Core interfaces receiving packets from other MPLS switches.	**ge-0/0/5.0** and **ge-0/0/6.0**	Interfaces that connect the provider switch to the ingress PE switch (PE-1). The EXP classifier is enabled by default on the switch and applied implicitly to these interfaces.
Core interfaces transmitting packets to other switches within the MPLS network.	**ge-0/0/7.0** and **ge-0/0/8.0**	Interfaces that transmit packets to the egress PE (PE-2). The EXP rewrite rule is applied implicitly on these interfaces. Schedulers can also be specified and will be applied to these interfaces.

Configuring the Local PE Switch

CLI Quick Configuration

To quickly configure a custom DSCP classifier, custom EXP rewrite rule, and a policer on the local PE switch, copy the following commands and paste them into the switch terminal window of PE-1:

```
[edit]
set class-of-service classifiers dscp dscp1 import default
set class-of-service classifiers dscp dscp1 forwarding-class expedited-forwarding loss-priority low code-points 000111
set class-of-service rewrite-rules exp e1 forwarding-class expedited-forwarding loss-priority low code-point 111
set class-of-service interfaces ge-0/0/1 unit 0 classifier dscp1
set firewall policer mypolicer if-exceeding bandwidth-limit 500m
```

```
set firewall policer mypolicer if-exceeding burst-size-limit 33553920
set firewall policer mypolicer then discard
set firewall family any filter myfilter term t1 then policer mypolicer
set protocols mpls label-switched-path lsp_to_pe2_ge1 to 127.1.1.3 policing filter myfilter
```

Step-by-Step Procedure

To configure a custom DSCP classifier, custom EXP rewrite rule, and a policer on the ingress PE switch:

1. Import the default DSCP classifier classes to the custom DSCP classifier that you are creating:

   ```
   [edit class-of-service]
   user@switch# set classifiers dscp dscp1 import default
   ```

2. Add the expedited-forwarding class to this custom DSCP classifier, specifying a loss priority and code point:

   ```
   [edit class-of-service]
   user@switch# set classifiers dscp dscp1 forwarding-class expedited-forwarding
   loss-priority low code-points 000111
   ```

3. Specify the values for the custom EXP rewrite rule, **e1**:

   ```
   [edit class-of-service]
   user@switch#  set rewrite-rules exp e1 forwarding-class expedited-forwarding
   loss-priority low code-point 111
   ```

4. Bind the DSCP classifier to the CCC interface:

   ```
   [edit ]
   user@switch#  set class-of-service interfaces ge-0/0/1 unit 0 classifier dscp1
   ```

5. Specify the number of bits per second permitted, on average, for the firewall policer, which will later be applied to the LSP:

   ```
   [edit firewall]
   set policer mypolicer if-exceeding bandwidth-limit 500m
   ```

6. Specify the maximum size permitted for bursts of data that exceed the given bandwidth limit for this policer:

   ```
   [edit firewall policer]
   set mypolicer if-exceeding burst-size-limit 33553920
   ```

7. Discard traffic that exceeds the rate limits for this policer:

   ```
   [edit firewall policer]
   set mypolicer then discard
   ```

8. To reference the policer, configure a filter term that includes the policer action:

   ```
   [edit firewall]
   user@switch#  set family any filter myfilter term t1 then policer mypolicer
   ```

9. Apply the filter to the LSP:

   ```
   [edit protocols mpls]
   set label-switched-path lsp_to_pe2_ge1 policing filter myfilter
   ```

Results

Display the results of the configuration:

```
[edit]
```

```
user@switch# show
class-of-service {
  classifiers {
    dscp dscp1 {
      import default;
      forwarding-class expedited-forwarding {
        loss-priority low code-points 000111;
      }
    }
  }
  interfaces {
    ge-0/0/1 {
      unit 0 {
        classifiers {
          dscp dscp1;
        }
      }
    }
  }
rewrite-rules {
  exp e1 {
    forwarding-class expedited-forwarding {
      loss-priority low code-point 111;
    }
  }
  }
}
firewall {
  family any {
    filter myfilter {
      term t1 {
        then policer mypolicer;
      }
    }
  }
  policer mypolicer {
    if-exceeding {
      bandwidth-limit 500m;
      burst-size-limit 33553920;
    }
    then discard;
  }
}
```

Configuring the Remote PE Switch

CLI Quick Configuration

To quickly configure a custom EXP classifier on the remote PE switch, copy the following commands and paste them into the switch terminal window of PE-2:

```
[edit]
set class-of-service classifiers exp exp1 import default
set class-of-service classifiers exp exp1 forwarding-class expedited-forwarding loss-priority low
code-points 010
```

Step-by-Step Procedure

To configure a custom EXP classifier on the egress PE switch:

1. Import the default EXP classifier classes to the custom EXP classifier that you are creating:

```
[edit class-of-service]
user@switch# set classifiers exp exp1 import default
```

2. Add the expedited-forwarding class to this custom EXP classifier, specifying a loss priority and code point:

```
[edit class-of-service]
user@switch# set classifiers exp exp1 forwarding-class expedited-forwarding loss-priority
low code-points 010
```

Results

Display the results of the configuration:

```
[edit]
 user@switch# show
class-of-service {
  classifiers {
    exp exp1 {
      import default;
      forwarding-class expedited-forwarding {
        loss-priority low code-points 010;
      }
    }
```

Configuring the Provider Switch

CLI Quick Configuration

To quickly configure a custom EXP classifier and a custom EXP rewrite rule on the provider switch, copy the following commands and paste them into the switch terminal window of the provider switch:

```
[edit]
set class-of-service classifiers exp exp1 import default
set class-of-service classifiers exp exp1 forwarding-class expedited-forwarding loss-priority low
code-points 010
set class-of-service rewrite-rules exp e1 forwarding-class expedited-forwarding loss-priority low
code-point 111
```

Step-by-Step Procedure

To configure a custom EXP classifier and a custom EXP rewrite rule on the provider switch:

1. Import the default EXP classifier classes to the custom EXP classifier that you are creating:

```
[edit class-of-service]
user@switch# set classifiers exp exp1 import default
```

2. Add the expedited-forwarding class to this custom EXP classifier, specifying a loss priority and code point:

```
[edit class-of-service]
user@switch# set classifiers exp exp1 forwarding-class expedited-forwarding loss-priority
low code-points 010
```

3. Specify the values for the custom EXP rewrite rule, **e1**:

```
[edit class-of-service]
user@switch# set rewrite-rules exp e1 forwarding-class expedited-forwarding
loss-priority low code-point 111
```

Results Display the results of the configuration:

```
[edit]
 user@switch# show
class-of-service {
  classifiers {
    exp exp1 {
      import default;
      forwarding-class expedited-forwarding {
        loss-priority low code-points 010;
      }
    }
  }
  rewrite-rules {
    exp e1 {
      forwarding-class expedited-forwarding {
        loss-priority low code-point 111;
      }
    }
  }
}
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying That the Policer Firewall Filter Is Operational on page 2905
- Verifying That the CoS Classifiers Are Going to the Right Queue on page 2905
- Verifying the CoS Forwarding Table Mapping on page 2908
- Verifying the Rewrite Rules on page 2909

Verifying That the Policer Firewall Filter Is Operational

Purpose Verify the operational state of the policer that is configured on the ingress PE switch.

Action

```
user@switch> show firewall
Filter: myfilter
Policers:
Name                                                Packets
mypolicer-t1                                              0
```

Meaning This output shows that the firewall filter **mypolicer** has been created.

Verifying That the CoS Classifiers Are Going to the Right Queue

Purpose Verify that the CoS classifiers are going to the right queue.

Action

```
user@switch> show class-of-service forwarding-table classifier

Classifier table index: 7, # entries: 64, Table type: DSCP
```

Entry #	Code point	Forwarding-class #	PLP
0	000000	0	0
1	000001	0	0
2	000010	0	0
3	000011	0	0
4	000100	0	0
5	000101	0	0
6	000110	0	0
7	000111	0	0
8	001000	0	0
9	001001	0	0
10	001010	0	0
11	001011	0	0
12	001100	0	0
13	001101	0	0
14	001110	0	0
15	001111	0	0
16	010000	0	0
17	010001	0	0
18	010010	0	0
19	010011	0	0
20	010100	0	0
21	010101	0	0
22	010110	0	0
23	010111	0	0
24	011000	0	0
25	011001	0	0
26	011010	0	0
27	011011	0	0
28	011100	0	0
29	011101	0	0
30	011110	0	0
31	011111	0	0
32	100000	0	0
33	100001	0	0
34	100010	0	0
35	100011	0	0
36	100100	0	0
37	100101	0	0
38	100110	0	0
39	100111	0	0
40	101000	0	0
41	101001	0	0
42	101010	0	0
43	101011	0	0
44	101100	0	0
45	101101	0	0
46	101110	0	0
47	101111	0	0
48	110000	3	0
49	110001	3	0
50	110010	3	0
51	110011	3	0
52	110100	3	0
53	110101	3	0
54	110110	3	0
55	110111	3	0
56	111000	3	0
57	111001	3	0
58	111010	3	0
59	111011	3	0

```
  60        111100        3       0
  61        111101        3       0
  62        111110        3       0
  63        111111        3       0

Classifier table index: 11, # entries: 8, Table type: IEEE 802.1
Entry #   Code point   Forwarding-class #   PLP
   0           000        0       0
   1           001        0       0
   2           010        0       0
   3           011        0       0
   4           100        0       0
   5           101        0       0
   6           110        3       0
   7           111        3       0

Classifier table index: 12, # entries: 8, Table type: IPv4 precedence
Entry #   Code point   Forwarding-class #   PLP
   0           000        0       0
   1           001        0       0
   2           010        0       0
   3           011        0       0
   4           100        0       0
   5           101        0       0
   6           110        3       0
   7           111        3       0

Classifier table index: 16, # entries: 8, Table type: Untrust
Entry #   Code point   Forwarding-class #   PLP
   0           000        0       0
   1           001        0       0
   2           010        0       0
   3           011        0       0
   4           100        0       0
   5           101        0       0
   6           110        0       0
   7           111        0       0

Classifier table index: 9346, # entries: 64, Table type: DSCP
Entry #   Code point   Forwarding-class #   PLP
   0        000000        0       0
   1        000001        0       0
   2        000010        0       0
   3        000011        0       0
   4        000100        0       0
   5        000101        0       0
   6        000110        0       0
   7        000111        1       0
   8        001000        0       0
   9        001001        0       0
  10        001010        0       0
  11        001011        0       0
  12        001100        0       0
  13        001101        0       0
  14        001110        0       0
  15        001111        0       0
  16        010000        0       0
  17        010001        0       0
  18        010010        0       0
  19        010011        0       0
  20        010100        0       0
```

21	010101	0	0
22	010110	0	0
23	010111	0	0
24	011000	0	0
25	011001	0	0
26	011010	0	0
27	011011	0	0
28	011100	0	0
29	011101	0	0
30	011110	0	0
31	011111	0	0
32	100000	0	0
33	100001	0	0
34	100010	0	0
35	100011	0	0
36	100100	0	0
37	100101	0	0
38	100110	0	0
39	100111	0	0
40	101000	0	0
41	101001	0	0
42	101010	0	0
43	101011	0	0
44	101100	0	0
45	101101	0	0
46	101110	0	0
47	101111	0	0
48	110000	3	0
49	110001	3	0
50	110010	3	0
51	110011	3	0
52	110100	3	0
53	110101	3	0
54	110110	3	0
55	110111	3	0
56	111000	3	0
57	111001	3	0
58	111010	3	0
59	111011	3	0
60	111100	3	0
61	111101	3	0
62	111110	3	0
63	111111	3	0

Meaning This output shows that a new DSCP classifier has been created, index **9346**, on the ingress PE switch (PE-1).

Verifying the CoS Forwarding Table Mapping

Purpose For each logical interface, display either the table index of the classifier for a given code point type or the queue number (if it is a fixed classification) in the forwarding table.

Action

```
user@switch>show class-of-service forwarding-table classifier mapping

                                Table Index/
```

```
Interface       Index      Q num      Table type
ge-0/0/1.0        92        9346      DSCP
```

Meaning The results show that the new DSCP classifier, index number **9346**, is bound to interface **ge-0/0/1.0**.

Verifying the Rewrite Rules

Purpose Display mapping of the queue number and loss priority to code point value for each rewrite rule as it exists in the forwarding table.

Action

```
user@switch>show class-of-service forwarding-table rewrite-rule

Rewrite table index: 31, # entries: 4, Table type: DSCP
FC#     Low bits  State      High bits  State
0       000000  Enabled       000000  Enabled
1       101110  Enabled       101110  Enabled
2       001010  Enabled       001100  Enabled
3       110000  Enabled       111000  Enabled

Rewrite table index: 34, # entries: 4, Table type: IEEE 802.1
FC#     Low bits  State      High bits  State
0          000  Enabled          001  Enabled
1          010  Enabled          011  Enabled
2          100  Enabled          101  Enabled
3          110  Enabled          111  Enabled

Rewrite table index: 35, # entries: 4, Table type: IPv4 precedence
FC#     Low bits  State      High bits  State
0          000  Enabled          000  Enabled
1          101  Enabled          101  Enabled
2          001  Enabled          001  Enabled
3          110  Enabled          111  Enabled

Rewrite table index: 9281, # entries: 1, Table type: EXP
FC#     Low bits  State      High bits  State
1          111  Enabled          000  Disabled
```

Meaning This output shows that a new EXP classifier with the index number **9281** has been created.

Related Documentation

- Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect (CLI Procedure) on page 3111
- Configuring MPLS on Provider Edge Switches Using IP Over MPLS (CLI Procedure) on page 3107
- Understanding Using CoS with MPLS Networks on J-EX Series Switches on page 2880
- Monitoring CoS Forwarding Classes on page 2936

CHAPTER 109

Configuring CoS

Configuring CoS (J-Web Procedure)

The Class of Service Configuration pages allow you to configure the Junos OS CoS components. You can configure forwarding classes for transmitting packets, define which packets are placed into each output queue, and schedule the transmission service level for each queue. After defining the CoS components you must assign classifiers to the required physical and logical interfaces.

Using the Class of Service Configuration pages, you can configure various CoS components individually or in combination to define particular CoS services.

To configure CoS components :

1. In the J-Web interface, select **Configure>Class of Service**.
2. On the Class of Service Configuration page, select one of the following options depending on the CoS component that you want to define. Enter information into the pages as described in the respective table:
   - To define or edit CoS value aliases, select **CoS Value Aliases** .
   - To define or edit forwarding classes and assign queues, select **Forwarding Classes**.
   - To define or edit classifiers, select **Classifiers** .
   - To define or edit rewrite rules, select **Rewrite Rules**.
   - To define or edit schedulers, select **Schedulers**.
   - To define or edit virtual channel groups, select **Interface Associations**.
3. Click **Apply** after completing configuration on any Configuration page.

Related Documentation

- Defining CoS Classifiers (J-Web Procedure) on page 2916
- Defining CoS Code-Point Aliases (J-Web Procedure) on page 2912
- Defining CoS Forwarding Classes (J-Web Procedure) on page 2918
- Defining CoS Rewrite Rules (J-Web Procedure) on page 2926
- Defining CoS Schedulers (J-Web Procedure) on page 2920
- Assigning CoS Components to Interfaces (J-Web Procedure) on page 2928

Defining CoS Code-Point Aliases (J-Web Procedure)

You can use the J-Web interface to define CoS code-point aliases on a J-EX Series switch. By defining aliases you can assign meaningful names to a particular set of bit values and refer to them when configuring CoS components.

To define CoS code-point aliases:

1. Select **Configure** > **Class of Service** > **CoS Value Aliases**.

NOTE: After you make changes to the configuration in this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options** > **Commit**. See “Using the Commit Options to Commit Configuration Changes (J-Web Procedure)” on page 334 for details about all commit options.

2. Click one:

   - **Add**—Adds a code-point alias. Enter information into the code point alias page as described in Table 369 on page 2913.
   - **Edit**—Modifies an existing code-point alias. Enter information into the code point alias page as described in Table 369 on page 2913.
   - **Delete**—Deletes an existing code-point alias.

Table 369 on page 2913 describes the related fields.

Table 369: CoS Value Aliases Configuration Fields

Field	Function	Your Action
Code point name	Specifies the name for a code-point—for example, **af11** or **be**.	Enter a name.
Code point type	Specifies a code-point type. The code-point type can be DSCP or IP precedence.	Select a value.
Code point value bits	Specifies the CoS value for which an alias is defined. Changing this value alters the behavior of all classifiers that refer to this alias.	To specify a CoS value, type it in the appropriate format: • For DSCP CoS values, use the format xxxxxx, where x is 1 or 0—for example, **101110**. • For IP precedence CoS values, use the format xxx, where x is 1 or 0—for example, **111**.

Related Documentation

- Defining CoS Code-Point Aliases (CLI Procedure) on page 2914
- Monitoring CoS Value Aliases on page 2940
- Example: Configuring CoS on J-EX Series Switches on page 2883

Defining CoS Code-Point Aliases (CLI Procedure)

You can use code-point aliases to streamline the process of configuring CoS features on your J-EX Series switch. A code-point alias assigns a name to a pattern of code-point bits. You can use this name instead of the bit pattern when you configure other CoS components such as classifiers, drop-profile maps, and rewrite rules.

You can configure code-point aliases for the following CoS marker types:

- DSCP—Handles incoming IPv4 packets.
- IEEE 802.1p—Handles Layer 2 CoS.
- Inet precedence—Handles incoming IPv4 packets. IP precedence mapping requires only the higher order three bits of the DSCP field.

To configure a code-point alias for a specified CoS marker type (**dscp**), assign an alias (**my1**) to the code-point (**110001**):

```
[edit class-of-service code-point-aliases]
user@switch# set dscp my1 110001
```

Related Documentation

- Defining CoS Code-Point Aliases (J-Web Procedure) on page 2912
- Example: Configuring CoS on J-EX Series Switches on page 2883
- Monitoring CoS Value Aliases on page 2940
- Understanding CoS Code-Point Aliases on page 2864

Defining CoS Classifiers (CLI Procedure)

Packet classification associates incoming packets with a particular CoS servicing level. Classifiers associate packets with a forwarding class and loss priority and assign packets to output queues based on the associated forwarding class. The Junos OS supports two general types of classifiers:

- Behavior aggregate or CoS value traffic classifiers—Examines the CoS value in the packet header. The value in this single field determines the CoS settings applied to the packet. BA classifiers allow you to set the forwarding class and loss priority of a packet based on the Differentiated Services code point (DSCP) value, IP precedence value, or IEEE 802.1p value.
- Multifield traffic classifiers—Examines multiple fields in the packet such as source and destination addresses and source and destination port numbers of the packet. With multifield classifiers, you set the forwarding class and loss priority of a packet based on firewall filter rules.

The following example describes how to configure a BA classifier **ba-classifier** as the default DSCP map and apply it to either a specific Gigabit Ethernet interface or to all the Gigabit Ethernet interfaces on the switch. The BA classifier assigns loss priorities, as shown in Table 370 on page 2915, to incoming packets in the four forwarding classes.

You can use the same procedure to set multifield classifiers (except that you would use firewall filter rules).

Table 370: BA-classifier Loss Priority Assignments

Forwarding Class	For CoS Traffic Type	ba-classifier Assignment
be	Best-effort traffic	High-priority code point: **000001**
ef	Expedited-forwarding traffic	High-priority code point: **101110**
af	Assured-forwarding traffic	High-priority code point: **001100**
nc	Network-control traffic	High-priority code point: **110001**

To configure a DSCP BA classifier named **ba-classifier** as the default DSCP map:

- Associate code point **000001** with forwarding class **be** and loss priority **high:**

```
[edit class-of-service classifiers]
user@switch# set dscp ba-classifier import default forwarding-class be loss-priority high code-points 000001
```

- Associate code point **101110** with forwarding class **ef** and loss priority **high:**

```
[edit class-of-service classifiers]
user@switch# set dscp ba-classifier forwarding-class ef loss-priority high code-points 101110
```

- Associate code point **001100** with forwarding class **af** and loss priority **high:**

```
[edit class-of-service classifiers]
user@switch# set dscp ba-classifier forwarding-class af loss-priority high code-points 001100
```

- Associate code point **110001** with forwarding class **nc** and loss priority **high:**

```
[edit class-of-service classifiers]
user@switch# set dscp ba-classifier forwarding-class nc loss-priority high code-points 110001
```

- Apply the classifier to a specific interface or to all Gigabit Ethernet interfaces on the switch.
  - To apply the classifier to a specific interface:

```
[edit class-of-service interfaces]
user@switch# set ge-0/0/0 unit 0 classifiers dscp ba-classifier
```

  - To apply the classifier to all Gigabit Ethernet interfaces on the switch, use wildcards for the interface name and the logical-interface (unit) number:

```
[edit class-of-service interfaces]
user@switch# set ge-* unit * classifiers dscp ba-classifier
```

Related Documentation

- Defining CoS Classifiers (J-Web Procedure) on page 2916
- Example: Configuring CoS on J-EX Series Switches on page 2883
- Assigning CoS Components to Interfaces (CLI Procedure) on page 2928
- Monitoring CoS Classifiers on page 2935
- Understanding CoS Classifiers on page 2867

Defining CoS Classifiers (J-Web Procedure)

You can use the J-Web interface to define CoS classifiers on a J-EX Series switch. Classifiers examine the CoS value or alias of an incoming packet and assign the packet a level of service by setting its forwarding class and loss priority.

To define CoS classifiers:

1. Select **Configure** > **Class of Service** > **Classifiers**.

   NOTE: After you make changes to the configuration in this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See "Using the Commit Options to Commit Configuration Changes (J-Web Procedure)" on page 334 for details about all commit options.

2. Click one:
   - **Add**—Adds a classifier. Enter information into the classifier page as described in Table 371 on page 2916.
   - **Edit**—Modifies an existing classifier. Enter information into the classifier page as described in Table 371 on page 2916.
   - **Delete**—Deletes an existing classifier.

Table 371: Classifiers Configuration Fields

Field	Function	Your Action
Classifier Name	Specifies the name for a classifier.	To name a classifier, type the name—for example, **ba-classifier**.
Classifier Type	Specifies the type of classifier: **dscp**, **ieee-802.1**, or **inet-precedence**.	Select a value from the list.

Table 371: Classifiers Configuration Fields *(continued)*

Field	Function	Your Action
Code Point Mapping	Sets the forwarding classes and the packet loss priorities (PLPs) for specific CoS values and aliases.	To add a code point mapping: 1. Click **Add**. 2. Select the code point. 3. Select a forwarding class from the following list: • **expedited-forwarding**—Provides low loss, low delay, low jitter, assured bandwidth, and end-to-end service. Packets can be forwarded out of sequence or dropped. • **best-effort**—Provides no special CoS handling of packets. Typically, RED drop profile is aggressive and no loss priority is defined. • **assured-forwarding**—Provides high assurance for packets within the specified service profile. Excess packets are dropped. • **network-control**—Packets can be delayed but not dropped. 4. Select the loss priority. To assign a loss priority, select one: • **high**—Packet has a high loss priority. • **low**—Packet has a low loss priority.

Related Documentation

- Defining CoS Classifiers (CLI Procedure) on page 2914
- Example: Configuring CoS on J-EX Series Switches on page 2883
- Monitoring CoS Classifiers on page 2935
- Understanding CoS Classifiers on page 2867

Defining CoS Forwarding Classes (CLI Procedure)

Forwarding classes allow you to group packets for transmission. Based on forwarding classes, you assign packets to output queues.

By default, four categories of forwarding classes are defined: best effort, assured forwarding, expedited forwarding, and network control. J-EX Series switches support up to 16 forwarding classes.

You can configure forwarding classes in one of the following ways:

- Using **class** statement—You can configure up to 16 forwarding classes and you can map multiple forwarding classes to single queue.
- Using **queue** statement—You can configure up to 8 forwarding classes and you can map one forwarding class to one queue.

  This example uses the **class** statement to configure forwarding classes.

To configure CoS forwarding classes, map the forwarding classes to queues:

```
[edit class-of-service forwarding-classes]
user@switch# set class be queue—num 0
user@switch# set class ef queue—num 1
user@switch# set class af queue—num 2
user@switch# set class nc queue—num 3
user@switch# set class ef1 queue—num 4
user@switch# set class ef2 queue—num 5
user@switch# set class af1 queue—num 6
user@switch# set class nc1 queue—num 7
```

Related Documentation

- Defining CoS Forwarding Classes (J-Web Procedure) on page 2918
- Example: Configuring CoS on J-EX Series Switches on page 2883
- Assigning CoS Components to Interfaces (CLI Procedure) on page 2928
- Monitoring CoS Forwarding Classes on page 2936
- Understanding CoS Forwarding Classes on page 2870

Defining CoS Forwarding Classes (J-Web Procedure)

You can define CoS forwarding classes on a J-EX Series switch using the J-Web interface. Assigning a forwarding class to a queue number affects the scheduling and marking of a packet as it transits a switch.

To define forwarding classes:

1. Select **Configure** > **Class of Service** > **Forwarding Classes**.

NOTE: After you make changes to the configuration in this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See "Using the Commit Options to Commit Configuration Changes (J-Web Procedure)" on page 334 for details about all commit options.

2. Click one:
   - **Add**—Adds a forwarding class. Enter information into the forwarding class page as described in Table 372 on page 2919.
   - **Edit**—Modifies an existing forwarding class. Enter information into the forwarding class page as described in Table 372 on page 2919.
   - **Delete**—Deletes an existing forwarding class.

Table 372: Forwarding Classes Configuration Fields

Field	Function	Your Action
Forwarding Class Summary		
Queue #	Specifies the internal queue numbers to which forwarding classes are assigned. By default, if a packet is not classified, it is assigned to the class associated with queue 0. You can have more than one forwarding class to a queue number.	To specify an internal queue number, select an integer from 0 through 7, appropriate for your platform.
Forwarding Class Name	Specifies the forwarding class names assigned to specific internal queue numbers. By default, four forwarding classes are assigned to queue numbers 0 (best-effort), 1 (assured-forwarding), 5 (expedited-forwarding), and 7 (network-connect).	Type the name—for example, **be-class**.

Related Documentation

- Defining CoS Forwarding Classes (CLI Procedure) on page 2918
- Example: Configuring CoS on J-EX Series Switches on page 2883
- Monitoring CoS Forwarding Classes on page 2936
- Assigning CoS Components to Interfaces (J-Web Procedure) on page 2928
- Understanding CoS Forwarding Classes on page 2870

Defining CoS Schedulers (CLI Procedure)

You use schedulers to define the CoS properties of output queues. These properties include the amount of interface bandwidth assigned to the queue, the size of the memory buffer allocated for storing packets, the priority of the queue, and the tail drop profiles associated with the queue.

You associate the schedulers with forwarding classes by means of scheduler maps. You can then associate each scheduler map with an interface, thereby configuring the queues and packet schedulers that operate according to this mapping.

You can associate up to four user-defined scheduler maps with the interfaces.

To configure CoS schedulers using the CLI:

1. Create a scheduler (**be-sched**) with low priority:

   ```
   [edit class-of-service schedulers]
   user@switch# set be-sched priority low
   ```

2. Configure a scheduler map (**be-map**) that associates the scheduler (**be-sched**) with the forwarding class (**best-effort**):

   ```
   [edit class-of-service scheduler-maps]
   user@switch# set be-map forwarding-class best-effort scheduler be-sched
   ```

3. Assign the scheduler map (**be-map**) to an Ethernet interface (**ge-0/0/1**):

   ```
   [edit class-of-service interfaces]
   user@switch# set ge-0/0/1 scheduler-map be-map
   ```

4. Alternatively to assign the scheduler map (**be-map**) to all the Ethernet interfaces using wild cards (**ge-***):

   ```
   [edit class-of-service interfaces]
   user@switch# set ge-* scheduler-map be-map
   ```

Related Documentation

- Defining CoS Schedulers (J-Web Procedure) on page 2920
- Example: Configuring CoS on J-EX Series Switches on page 2883
- Assigning CoS Components to Interfaces (CLI Procedure) on page 2928
- Monitoring CoS Scheduler Maps on page 2939
- Understanding CoS Schedulers on page 2873

Defining CoS Schedulers (J-Web Procedure)

You can use the J-Web interface to define CoS schedulers on a J-EX Series switch. Using schedulers, you can assign attributes to queues and thereby provide congestion control for a particular class of traffic. These attributes include the amount of interface bandwidth, memory buffer size, transmit rate, and schedule priority.

To configure schedulers:

1. Select **Configure** > **Class of Service** > **Schedulers**.

NOTE: After you make changes to the configuration in this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See "Using the Commit Options to Commit Configuration Changes (J-Web Procedure)" on page 334 for details about all commit options.

2. Click one:

   - **Add**—Adds a scheduler. Enter information into the schedulers page as described in Table 373 on page 2921.
   - **Edit**—Modifies an existing scheduler. Enter information into the schedulers page as described in Table 373 on page 2921.
   - **Delete**—Deletes an existing scheduler.

Table 373: Schedulers Configuration Page

Field	Function	Your Action
Scheduler Name	Specifies the name for a scheduler.	To name a scheduler, type the name—for example, **be-scheduler**.
Scheduling Priority	Sets the transmission priority of the scheduler, which determines the order in which an output interface transmits traffic from the queues. You can set scheduling priority at different levels in the order of increasing priority from low to high. A high-priority queue with a high transmission rate might lock out lower-priority traffic.	To set a priority, select one: • **low**—Packets in this queue are transmitted last. • **strict-high**—Packets in this queue are transmitted first. • To specify no scheduling priority, select the blank.

Table 373: Schedulers Configuration Page *(continued)*

Field	Function	Your Action
Buffer Size	Defines the size of the delay buffer. By default, queues 0 through 7 are allotted the following percentage of the total available buffer space: • Queue 0—95 percent • Queue 1—0 percent • Queue 2—0 percent • Queue 3—0 percent • Queue 4—0 percent • Queue 5—0 percent • Queue 6—0 percent • Queue 7—5 percent NOTE: A large buffer size value correlates with a greater possibility of packet delays. Such a value might not be practical for sensitive traffic such as voice or video.	To define a delay buffer size for a scheduler, select the appropriate option: • To specify no buffer size, select the blank. • To specify buffer size as a percentage of the total buffer, select **Percent** and type an integer from 1 through 100. • To specify buffer size as the remaining available buffer, select **Remainder**. NOTE: On J-EX8200 switches, you can specify the buffer size as a temporal value. The queuing algorithm will then drop packets once it has queued a computed number of bytes. This number is the product of the logical interface speed and the configured temporal value.
Shaping Rate	Specifies the rate at which queues transmit packets.	• To specify shaping rate as a percentage, select **Percent** and type an integer from 1 through 100. • To specify shaping rate as a number, select **Rate** and enter a value. • To specify no shaping rate, select the blank.
Transmit Rate	Defines the transmission rate of a scheduler. The transmit rate determines the traffic bandwidth from each forwarding class you configure. By default, queues 0 through 7 are allotted the following percentage of the transmission capacity: • Queue 0—95 percent • Queue 1—0 percent • Queue 2—0 percent • Queue 3—5 percent • Queue 4—0 percent • Queue 6—0 percent • Queue 7—5 percent	To define a transmit rate, select the appropriate option: • To enforce the exact transmission rate, select **Rate** and enter a value. • To specify the remaining transmission capacity, select **Remainder Available**. • To specify a percentage of transmission capacity, select **Percent** and type an integer from 1 through 100. • To specify no transmit rate, select the blank.

Related Documentation

- Defining CoS Schedulers (CLI Procedure) on page 2920
- Example: Configuring CoS on J-EX Series Switches on page 2883
- Monitoring CoS Scheduler Maps on page 2939

Defining CoS Scheduler Maps (J-Web Procedure)

You can use the J-Web interface to configure CoS scheduler maps on a J-EX Series switch.

To configure scheduler maps:

1. Select **Configure** > **Class of Service** > **Scheduler Maps**.

   NOTE: After you make changes to the configuration in this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options** > **Commit**. See "Using the Commit Options to Commit Configuration Changes (J-Web Procedure)" on page 334 for details about all commit options.

2. Click one:

   - **Add**—Adds a scheduler map. Enter information into the scheduler map page as described in Table 374 on page 2923.
   - **Edit**—Modifies an existing scheduler map. Enter information into the scheduler map page as described in Table 374 on page 2923.
   - **Delete**—Deletes an existing scheduler map.

Table 374: Scheduler Maps Configuration Fields

Field	Function	Your Action
Scheduler Map Name	Specifies the name for a scheduler map.	To name a map, type the name—for example, **be-scheduler-map**.
Scheduler Mapping	Allows you to associate a preconfigured scheduler with a forwarding class. After scheduler maps have been applied to an interface, they affect the hardware queues and packet schedulers.	To associate a scheduler with a forwarding class, locate the forwarding class and select the scheduler in the box next to it. For example, for the **best-effort** forwarding class, select the configured scheduler from the list.

Related Documentation

- Defining CoS Schedulers (J-Web Procedure) on page 2920
- Defining CoS Schedulers (CLI Procedure) on page 2920
- Example: Configuring CoS on J-EX Series Switches on page 2883
- Monitoring CoS Scheduler Maps on page 2939

Defining CoS Drop Profiles (J-Web Procedure)

You can use the J-Web interface to define CoS drop profiles on J-EX8200 switches.

To configure CoS drop profiles:

1. Select **Configure** > **Class of Service** > **Drop Profile**.

NOTE: After you make changes to the configuration in this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options** > **Commit**. See "Using the Commit Options to Commit Configuration Changes (J-Web Procedure)" on page 334 for details about all commit options.

2. Click one:
   - **Add**—Adds a drop profile. Enter information into the drop profiles page as described in Table 375 on page 2924.
   - **Edit**—Modifies an existing drop file. Enter information into the drop profiles page as described in Table 375 on page 2924.
   - **Delete**—Deletes an existing drop profile.

Table 375: Drop Profiles Configuration parameters

Field	Function	Your Action
Drop Profile Name	Specifies the name for a drop profile.	Type the name.
Drop profile graph	Specifies the drop profile graph type	Select one: **Segmented** or **Interpolated**.
Drop profile values	Specifies values for the following two parameters of the drop profile: the queue fill level and the drop probability. The queue fill level represents a percentage of the memory used to store packets in relation to the total amount that has been allocated for that specific queue. The drop probability is a percentage value that correlates to the likelihood that an individual packet is dropped from the network.	To add new values: 1. Click **Add**. 2. Enter the fill level. 3. Enter the drop probability. 4. Click **OK**. To edit an existing value, click **Edit** and modify the fill level and drop probability. To delete a value, select it and click **Delete**.

Related Documentation

- Monitoring CoS Drop Profiles on page 2941
- Example: Configuring CoS on J-EX Series Switches on page 2883

Configuring CoS Tail Drop Profiles (CLI Procedure)

Tail drop is a simple and effective traffic congestion avoidance mechanism. When you apply this mechanism to manage congestion, packets are dropped when the output queue is full.

To configure CoS tail-drop profiles, create a drop profile name (**be-dp**) and assign a fill level (**25**):

```
[edit class-of-service drop-profiles]
user@switch# set be-dp fill-level 25
```

Related Documentation

- Example: Configuring CoS on J-EX Series Switches on page 2883
- Understanding CoS Tail Drop Profiles on page 2872

Defining CoS Rewrite Rules (CLI Procedure)

You configure rewrite rules to alter CoS values in outgoing packets on the outbound interfaces of a J-EX Series switch to match the policies of a targeted peer. Policy matching allows the downstream routing platform or switch in a neighboring network to classify each packet into the appropriate service group.

To configure a CoS rewrite rule, create the rule by giving it a name and associating it with a forwarding class, loss priority, and a code point, thus creating a rewrite table. After the rewrite rule is created, enable it on an interface. You can also apply an existing rewrite rule on an interface.

NOTE: To replace an existing rewrite rule on the interface with a new rewrite rule of the same type, first explicitly remove the rewrite rule and then apply the new rule.

NOTE: Custom rewrite-rule bindings are implemented through filters. And custom rewrite rules cannot be bound to routed VLAN interfaces (RVIs).

To create rewrite rules and enable them on interfaces:

- To create an 802.1p rewrite rule named customup-rw in the rewrite table for all Layer 2 interfaces:

```
[edit class-of-service rewrite-rules]
user@switch# set ieee-802.1 customup-rw forwarding-class be loss-priority low
code-point 000
user@switch# set ieee-802.1 customup-rw forwarding-class be loss-priority high
code-point 001
user@switch# set ieee-802.1 customup-rw forwarding-class af loss-priority low
code-point 010
user@switch# set ieee-802.1 customup-rw forwarding-class af loss-priority high
code-point 011
```

```
user@switch# set ieee-802.1 customup-rw forwarding-class ef loss-priority low
code-point 100
user@switch# set ieee-802.1 customup-rw forwarding-class ef loss-priority high
code-point 101
user@switch# set ieee-802.1 customup-rw forwarding-class nc loss-priority low
code-point 110
user@switch# set ieee-802.1 customup-rw forwarding-class nc loss-priority high
code-point 111
```

- To enable an 802.1p rewrite rule named customup-rw on a Layer 2 interface:

```
[edit]
user@switch# set class-of-service interfaces ge-0/0/0 unit 0 rewrite-rules ieee-802.1
customup-rw
```

- To enable an 802.1p rewrite rule named customup-rw on all Gigabit Ethernet interfaces on the switch, use wildcards for the interface name and logical-interface (unit) number:

```
[edit]
user@switch# set class-of-service interfaces ge-* unit * rewrite-rules customup-rw
```

Related Documentation

- Defining CoS Rewrite Rules (J-Web Procedure) on page 2926
- Example: Configuring CoS on J-EX Series Switches on page 2883
- Monitoring CoS Rewrite Rules on page 2938
- Understanding CoS Rewrite Rules on page 2876

Defining CoS Rewrite Rules (J-Web Procedure)

You can use the J-Web interface to define CoS rewrite rules. Use the rewrite rules to alter the CoS values in outgoing packets to meet the requirements of the targeted peer. A rewrite rule examines the forwarding class and loss priority of a packet and sets its bits to a corresponding value specified in the rule.

To define rewrite rules:

1. Select **Configure > Class of Service > Rewrite Rules**.

NOTE: After you make changes to the configuration in this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See "Using the Commit Options to Commit Configuration Changes (J-Web Procedure)" on page 334 for details about all commit options.

2. Click one:
   - **Add**—Adds a rewrite rule. Enter information into the rewrite rule page as described in Table 376 on page 2927.
   - **Edit**—Modifies an existing rewrite rule. Enter information into the rewrite rule page as described in Table 376 on page 2927.

- **Delete**—Deletes an existing rewrite rule.

Table 376: Rewrite Rules Configuration Page Summary

Field	Function	Your Action
Rewrite Rule Name	Specifies the name for the rewrite rule.	To name a rule, type the name—for example, **rewrite-dscps**.
Rewrite rule type	Specifies the type of rewrite rule: **dscp**, **ieee-802.1**, or **inet-precedence**.	Select a value from the list.
Code Point Mapping	Rewrites outgoing CoS values of a packet based on the forwarding class and loss priority. Allows you to remove a code point mapping entry.	To configure a CoS value assignment, follow these steps: To add a code point mapping: 1. Click **Add**. 2. Select the code point. 3. Select a forwarding class from the following list: • **expedited-forwarding**—Provides low loss, low delay, low jitter, assured bandwidth, and end-to-end service. Packets can be forwarded out of sequence or dropped. • **best-effort**—Provides no special CoS handling of packets. Typically, RED drop profile is aggressive and no loss priority is defined. • **assured-forwarding**—Provides high assurance for packets within the specified service profile. Excess packets are dropped. • **network-control**—Packets can be delayed but not dropped. 4. Select the loss priority. To assign a loss priority, select one: • **high**—Packet has a high loss priority. • **low**—Packet has a low loss priority. To edit an existing code point mapping, select it and click **Edit**. To remove a code point mapping entry, select it and click **Remove**.

Related Documentation

- Defining CoS Rewrite Rules (CLI Procedure) on page 2925
- Understanding CoS Rewrite Rules on page 2876
- Monitoring CoS Rewrite Rules on page 2938
- Example: Configuring CoS on J-EX Series Switches on page 2883

Assigning CoS Components to Interfaces (CLI Procedure)

After you have defined the following CoS components, you must assign them to logical or physical interfaces.

- Forwarding classes—Assign only to logical interfaces.
- Classifiers—Assign only to logical interfaces.
- Scheduler maps—Assign to either physical or logical interfaces.
- Rewrite rules—Assign to either physical or logical interfaces.

You can assign a CoS component to a single interface or to multiple interfaces using wild cards.

To assign CoS components to interfaces:

- To assign CoS components to a single interface, associate a CoS component (for example a scheduler map named **ethernet-cos-map**) with an interface:

```
[edit class-of-service interfaces]
user@switch# set ge-0/0/20 scheduler-map ethernet-cos-map
```

- To assign a CoS component to multiple interfaces, associate a CoS component (for example, a rewrite rule named **customup-rw**) to all Gigabit Ethernet interfaces on the switch, use wild characters for the interface name and logical-interface (unit) number:

```
[edit class-of-service interfaces]
user@switch# set ge-* unit * rewrite-rules ieee-802.1 customup-rw
```

Related Documentation

- Assigning CoS Components to Interfaces (J-Web Procedure) on page 2928
- Example: Configuring CoS on J-EX Series Switches on page 2883
- Monitoring Interfaces That Have CoS Components on page 2937
- Understanding Junos OS CoS Components for J-EX Series Switches on page 2862

Assigning CoS Components to Interfaces (J-Web Procedure)

After you have defined CoS components on a J-EX Series switch, you must assign them to logical or physical interfaces. You can use the J-Web interface to assign scheduler maps to physical or logical interfaces and to assign forwarding classes or classifiers to logical interfaces.

To assign CoS components to interfaces:

1. Select **Configure** > **Class of Service** > **Assign to Interface**.

   NOTE: After you make changes to the configuration in this page, you must commit the changes immediately for them to take effect. To commit all changes to the active configuration, select **Commit Options > Commit**. See "Using the Commit Options to Commit Configuration Changes (J-Web Procedure)" on page 334 for details about all commit options.

2. To configure interface association, select an interface from the list and click **Edit**.
3. Select one:
   - **Associate system default scheduler map**—Associates the interface with the default scheduler map.
   - **Select the scheduler map**—Associates the interface with a configured scheduler map. Select the scheduler map from the list.
4. Click **OK**.
5. To manage a CoS service assignment on a logical interface, click one:
   - **Add**—Adds a CoS service to a logical interface on a specified physical interface. Enter information as described in Table 377 on page 2929.
   - **Edit**—Modifies a CoS service assignment to a logical interface. Enter information as described in Table 377 on page 2929.
   - **Delete**—Deletes the CoS service assignment to a logical interface.

Table 377: Assigning CoS Components to Logical Interfaces

Field	Function	Your Action
Unit	Specifies the name of a logical interface. Allows you to assign CoS components while configuring a logical interface on a physical interface at the same time.	Type the interface name. To assign CoS services to all logical interfaces configured on this physical interface, type the wildcard character (*).
Forwarding Class	Assigns a predefined forwarding class to incoming packets on a logical interface.	To assign a forwarding class to an interface, select the forwarding class.
Classifiers	Allows you to apply classification maps to a logical interface. Classifiers assign a forwarding class and loss priority to an incoming packet based on its CoS value.	To assign a classification map to an interface, select an appropriate classifier for each CoS value type used on the interface.
Rewrite Rules	Allows you to alter the CoS values in outgoing packets to meet the requirements of the targeted peer. A rewrite rule examines the forwarding class and loss priority of a packet and sets its bits to a corresponding value specified in the rule.	To assign rewrite rules to the interface, select the appropriate rewrite rule for each CoS value type used on the interface.

Related Documentation

- Assigning CoS Components to Interfaces (CLI Procedure) on page 2928
- Example: Configuring CoS on J-EX Series Switches on page 2883
- Monitoring Interfaces That Have CoS Components on page 2937

Configuring Junos OS EZQoS for CoS (CLI Procedure)

You use Junos OS EZQoS on J-EX Series switches to eliminate the complexities involved in configuring class of service (CoS) across the network. EZQoS offers templates for key traffic classes.

When you configure EZQoS on J-EX Series switches, preconfigured values are assigned to all CoS parameters based on the typical application requirements. These preconfigured values are stored in a template with a unique name.

NOTE: Currently, we provide an EZQoS template for configuring CoS for VoIP applications. The EZQoS VoIP template is stored in /etc/config/ezqos-voip.conf.

To configure EZQoS using the CLI:

1. Load the EZQoS configuration file (**/etc/config/ezqos-voip.conf**):

   ```
   [edit]
   user@switch# load merge /etc/config/ezqos-voip.conf
   ```

2. Apply the EZQoS group (**ezqos-voip**):

   ```
   [edit]
   user@switch# set apply-groups ezqos-voip
   ```

3. Apply the DSCP classifier (**ezqos-dscp-classifier**) to a Gigabit Ethernet interface (**ge-0/0/0**):

   ```
   [edit class-of-service interfaces]
   user@switch# set ge-0/0/0 unit 0 classifiers dscp ezqos-dscp-classifier
   ```

4. Apply the scheduler map (**ezqos-voip-sched-maps**) to a Gigabit Ethernet interface (**ge-0/0/1**):

   ```
   [edit class-of-service interfaces]
   user@switch# set ge-0/0/1 scheduler-map ezqos-voip-sched-maps
   ```

Related Documentation

- Example: Configuring CoS on J-EX Series Switches on page 2883
- Understanding Junos OS EZQoS for CoS Configurations on J-EX Series Switches on page 2879

Configuring CoS on MPLS Provider Edge Switch Using IP Over MPLS (CLI Procedure)

You can use class of service (CoS) within MPLS networks to prioritize certain types of traffic during periods of congestion. This topic describes configuring CoS components on a provider edge (PE) switch that is using IP Over MPLS.

This task describes now to create a custom DSCP classifier and a custom EXP rewrite rule on the ingress PE switch. It includes configuring a policer firewall filter and applying it to the customer-edge interface of the ingress PE switch. The policer firewall filter ensures that the amount of traffic forwarded through the MPLS tunnel never exceeds the requested bandwidth allocation.

For this procedure, we assume that the switch has already been configured for MPLS. See "Configuring MPLS on Provider Edge Switches Using MPLS Over IP (CLI Procedure)" on page 3107.

1. Import the default DSCP classifier classes to the custom DSCP classifier that you are creating:

   ```
   [edit class-of-service]
   user@switch#set classifiers dscp dscp1 import default
   ```

2. Add the expedited-forwarding class to this custom DSCP classifier, specifying a loss priority and code point:

   ```
   [edit class-of-service]
   user@switch#set classifiers dscp dscp1 forwarding-class expedited-forwarding
   loss-priority low code-points 000111
   ```

3. Specify the values for the custom EXP rewrite rule, **e1**:

   ```
   [edit class-of-service]
   user@switch#  set rewrite-rules exp e1 forwarding-class expedited-forwarding
   loss-priority low code-point 111
   ```

4. Specify the number of bits per second permitted, on average, for the firewall policer, which will later be applied to the customer-edge-interface:

   ```
   [edit firewall]
   set policer mypolicer if-exceeding bandwidth-limit 500m
   ```

5. Specify the maximum size permitted for bursts of data that exceed the given bandwidth limit for this policer:

   ```
   [edit firewall policer]
   set mypolicer if-exceeding burst-size-limit 33553920
   ```

6. Discard traffic that exceeds the rate limits for this policer:

   ```
   [edit firewall policer]
   set mypolicer then discard
   ```

7. To reference the policer, configure a filter term that includes the policer action:

   ```
   [edit firewall]
   user@switch#  set family inet filter myfilter term t1 then policer mypolicer
   ```

8. Apply the filter to the customer-edge interface:

```
[edit interfaces]
user@switch# set ge-2/0/3 unit 0 family inet address 121.121.121.1/16 policing filter
myfilter
```

NOTE: You can also configure schedulers and shapers as needed. See “Defining CoS Schedulers (CLI Procedure)” on page 2920.

Related Documentation

- Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect (CLI Procedure) on page 3111
- Assigning CoS Components to Interfaces (CLI Procedure) on page 2928
- Configuring Policers to Control Traffic Rates (CLI Procedure) on page 2788
- Understanding the Use of Policers in Firewall Filters on page 2752

Configuring CoS on MPLS Provider Edge Switch Using Circuit Cross-Connect (CLI Procedure)

You can use class of service (CoS) within MPLS networks to prioritize certain types of traffic during periods of congestion. This topic describes configuring CoS components on provider edge (PE) switch that is using MPLS over circuit-cross connect (CCC).

NOTE: If you are using MPLS with CCC, you can use only one type of DSCP/IP precedence and only one type of IEEE 802.1p on the CCC interfaces.

This procedure creates a custom DSCP classifier and a custom EXP rewrite rule on the ingress PE. It also enables a policer on the label-switched path (LSP) of the ingress PE to ensure that the amount of traffic forwarded through the LSP never exceeds the requested bandwidth allocation.

1. Import the default DSCP classifier classes to the custom DSCP classifier that you are creating:

   ```
   [edit class-of-service]
   user@switch#set classifiers dscp dscp1 import default
   ```

2. Add the expedited-forwarding class to this custom DSCP classifier, specifying a loss priority and code point:

   ```
   [edit class-of-service]
   user@switch#set classifiers dscp dscp1 forwarding-class expedited-forwarding
   loss-priority low code-points 000111
   ```

3. Specify the values for the custom EXP rewrite rule, **e1**:

   ```
   [edit class-of-service]
   user@switch# set rewrite-rules exp e1 forwarding-class expedited-forwarding
   loss-priority low code-point 111
   ```

4. Bind the DSCP classifier to the CCC interface:

```
[edit ]
user@switch# set class-of-service interfaces ge-0/0/1 unit 0 classifier dscp1
```

5. Specify the number of bits per second permitted, on average, for the firewall policer, which will later be applied to the LSP:

```
[edit firewall]
set policer mypolicer if-exceeding bandwidth-limit 500m
```

6. Specify the maximum size permitted for bursts of data that exceed the given bandwidth limit for this policer:

```
[edit firewall policer]
set mypolicer if-exceeding burst-size-limit 33553920
```

7. Discard traffic that exceeds the rate limits for this policer:

```
[edit firewall policer]
set mypolicer then discard
```

8. To reference the policer, configure a filter term that includes the policer action:

```
[edit firewall]
user@switch# set family any filter myfilter term t1 then policer mypolicer
```

9. Apply the filter to the LSP:

```
[edit protocols mpls]
set label-switched-path lsp_to_pe2_ge1 policing filter myfilter
```

NOTE: You can also configure schedulers and shapers as needed. See “Defining CoS Schedulers (CLI Procedure)” on page 2920.

Related Documentation

- Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect (CLI Procedure) on page 3111
- Assigning CoS Components to Interfaces (CLI Procedure) on page 2928
- Configuring Policers to Control Traffic Rates (CLI Procedure) on page 2788
- Understanding the Use of Policers in Firewall Filters on page 2752

CHAPTER 110

Verifying CoS Configuration

Monitoring CoS Classifiers

Purpose Use the monitoring functionality to display the mapping of incoming CoS values to forwarding class and loss priority for each classifier.

Action To monitor CoS classifiers in the J-Web interface, select **Monitor>Class of Service>Classifiers**

To monitor CoS classifiers in the CLI, enter the following CLI command:

show class-of-service classifier

Meaning Table 378 on page 2935 summarizes key output fields for CoS classifiers.

Table 378: Summary of Key CoS Classifier Output Fields

Field	Values	Additional Information
Classifier Name	Name of a classifier.	To display classifier assignments, click the plus sign (+).
CoS Value Type	The classifiers are displayed by type: • **dscp**—All classifiers of the DSCP type. • **ieee-802.1**—All classifiers of the IEEE 802.1 type. • **inet-precedence**—All classifiers of the IP precedence type.	
Index	Internal index of the classifier.	

Table 378: Summary of Key CoS Classifier Output Fields *(continued)*

Field	Values	Additional Information
Incoming CoS Value	CoS value of the incoming packets, in bits. These values are used for classification.	
Assign to Forwarding Class	Forwarding class that the classifier assigns to an incoming packet. This class affects the forwarding and scheduling policies that are applied to the packet as it transits the switch.	
Assign to Loss Priority	Loss priority value that the classifier assigns to the incoming packet based on its CoS value.	

Related Documentation

- Defining CoS Classifiers (CLI Procedure) on page 2914
- Defining CoS Classifiers (J-Web Procedure) on page 2916
- Example: Configuring CoS on J-EX Series Switches on page 2883

Monitoring CoS Forwarding Classes

Purpose View the current assignment of class-of-service (CoS) forwarding classes to queues on the switch.

Action To monitor CoS forwarding classes in the J-Web interface, select **Monitor>Class of Service>Forwarding Classes**.

To monitor CoS forwarding classes in the CLI, enter the following CLI command:

show class-of-service forwarding-class

Meaning Table 379 on page 2936 summarizes key output fields for CoS forwarding classes.

Table 379: Summary of Key CoS Forwarding Class Output Fields

Field	Values
Forwarding Class	Names of forwarding classes assigned to queue numbers. By default, the following forwarding classes are assigned to queues 0, 1, 5, or 7: • **best-effort**—Provides no special CoS handling of packets. Loss priority is typically not carried in a CoS value. • **expedited-forwarding**—Provides low loss, low delay, low jitter, assured bandwidth, and end-to-end service. • **assured-forwarding**—Provides high assurance for packets within specified service profile. Excess packets are dropped. • **network-control**—Packets can be delayed but not dropped.

Table 379: Summary of Key CoS Forwarding Class Output Fields *(continued)*

Field	Values
Queue	Queue number corresponding to the forwarding class name. By default, four queues, 0, 1, 5, or 7, are assigned to forwarding classes.

Related Documentation

- Defining CoS Forwarding Classes (CLI Procedure) on page 2918
- Defining CoS Forwarding Classes (J-Web Procedure) on page 2918
- Example: Configuring CoS on J-EX Series Switches on page 2883

Monitoring Interfaces That Have CoS Components

Purpose Use the monitoring functionality to display details about the physical and logical interfaces and the CoS components assigned to them.

Action To monitor interfaces that have CoS components in the J-Web interface, select **Monitor>Class of Service>Interface Association**.

To monitor interfaces that have CoS components in the CLI, enter the following command:

show class-of-service interface *interface*

Meaning Table 380 on page 2937 summarizes key output fields for CoS interfaces.

Table 380: Summary of Key CoS Interfaces Output Fields

Field	Values	Additional Information
Interface	Name of a physical interface to which CoS components are assigned.	To display names of logical interfaces configured on this physical interface, click the plus sign (+).
Scheduler Map	Name of the scheduler map associated with this interface.	
Queues Supported	Number of queues you can configure on the interface.	
Queues in Use	Number of queues currently configured.	
Logical Interface	Name of a logical interface on the physical interface to which CoS components are assigned.	
Object	Category of an object—for example, **classifier**, **scheduler-map**, or **rewrite**.	
Name	Name that you have given to an object—for example, **ba-classifier**.	

Table 380: Summary of Key CoS Interfaces Output Fields *(continued)*

Field	Values	Additional Information
Type	Type of an object—for example, **dscp** for a classifier.	
Index	Index of this interface or the internal index of a specific object.	

Related Documentation

- Assigning CoS Components to Interfaces (CLI Procedure) on page 2928
- Assigning CoS Components to Interfaces (J-Web Procedure) on page 2928
- Example: Configuring CoS on J-EX Series Switches on page 2883

Monitoring CoS Rewrite Rules

Purpose Use the monitoring functionality to display information about CoS value rewrite rules, which are based on the forwarding class and loss priority.

Action To monitor CoS rewrite rules in the J-Web interface, select **Monitor>Class of Service>Rewrite Rules**.

To monitor CoS rewrite rules in the CLI, enter the following command:

show class-of-service rewrite-rules

Meaning Table 381 on page 2938 summarizes key output fields for CoS rewrite rules.

Table 381: Summary of Key CoS Rewrite Rules Output Fields

Field	Values	Additional Information
Rewrite Rule Name	Names of rewrite rules.	
CoS Value Type	Rewrite rule type: • **dscp**—For IPv4 DiffServ traffic. • **exp**—For MPLS traffic. • **ieee-802.1**—For Layer 2 traffic. • **inet-precedence**—For IPv4 traffic.	To display forwarding classes, loss priorities, and rewritten CoS values, click the plus sign (+).
Index	Internal index for this particular rewrite rule.	
Forwarding Class	Forwarding class that is used to determine CoS values for rewriting in combination with loss priority.	Rewrite rules are applied to CoS values in outgoing packets based on forwarding class and loss priority setting.
Loss Priority	Loss priority that is used to determine CoS values for rewriting in combination with forwarding class.	

Table 381: Summary of Key CoS Rewrite Rules Output Fields *(continued)*

Field	Values	Additional Information
Rewrite CoS Value To	Value that the CoS value is rewritten to.	

Related Documentation

- Defining CoS Rewrite Rules (CLI Procedure) on page 2925
- Defining CoS Rewrite Rules (J-Web Procedure) on page 2926
- Example: Configuring CoS on J-EX Series Switches on page 2883

Monitoring CoS Scheduler Maps

Purpose Use the monitoring functionality to display assignments of CoS forwarding classes to schedulers.

Action To monitor CoS scheduler maps in the J-Web interface, select **Monitor>Class of Service>Scheduler Maps**.

To monitor CoS scheduler maps in the CLI, enter the following CLI command:

show class-of-service scheduler-map

Meaning Table 382 on page 2939 summarizes key output fields for CoS scheduler maps.

Table 382: Summary of Key CoS Scheduler Maps Output Fields

Field	Values	Additional Information
Scheduler Map	Name of a scheduler map.	For details, click the plus sign (+).
Index	Index of a specific object—scheduler maps, schedulers, or drop profiles.	
Scheduler Name	Name of a scheduler.	
Forwarding Class	Forwarding classes this scheduler is assigned to.	
Transmit Rate	Configured transmit rate of the scheduler in bits per second (bps). The rate value can be either of the following: • A percentage—The scheduler receives the specified percentage of the total interface bandwidth. • **remainder**— The scheduler receives the remaining bandwidth of the interface after bandwidth allocation to other schedulers.	

Table 382: Summary of Key CoS Scheduler Maps Output Fields *(continued)*

Field	Values	Additional Information
Buffer Size	Delay buffer size in the queue or the amount of transmit delay (in milliseconds). The buffer size can be either of the following: • A percentage—The buffer is a percentage of the total buffer allocation. • **remainder**—The buffer is sized according to what remains after other scheduler buffer allocations.	
Priority	Scheduling priority of a queue: • **strict-high**—Packets in this queue are transmitted first. • **low**—Packets in this queue are transmitted last.	
Drop Profiles	Name and index of a drop profile that is assigned to a specific loss priority and protocol pair.	
Loss Priority	Packet loss priority corresponding to a drop profile.	
Protocol	Transport protocol corresponding to a drop profile.	
Drop Profile Name	Name of the drop profile.	
Index	Index of a specific object—scheduler maps, schedulers, or drop profiles.	

Related Documentation

- Defining CoS Schedulers (CLI Procedure) on page 2920
- Defining CoS Schedulers (J-Web Procedure) on page 2920
- Example: Configuring CoS on J-EX Series Switches on page 2883

Monitoring CoS Value Aliases

Purpose Use the monitoring functionality to display information about the CoS value aliases that the system is currently using to represent DSCP, IEEE 802.1p, and IPv4 precedence bits.

Action To monitor CoS value aliases in the J-Web interface, select **Monitor>Class of Service>CoS Value Aliases**.

To monitor CoS value aliases in the CLI, enter the following command:

show class-of-service code-point-aliases

Meaning Table 383 on page 2941 summarizes key output fields for CoS value aliases.

Table 383: Summary of Key CoS Value Alias Output Fields

Field	Values	Additional Information
CoS Value Type	Type of the CoS value: • **dscp**—Examines Layer 3 packet headers for IP packet classification. • **ieee-802.1**—Examines Layer 2 packet headers for packet classification. • **inet-precedence**—Examines Layer 3 packet headers for IP packet classification.	To display aliases and bit patterns, click the plus sign (+).
CoS Value Alias	Name given to a set of bits—for example, **af11** is a name for **001010** bits.	
CoS Value	Set of bits associated with an alias.	

Related Documentation
- Defining CoS Code-Point Aliases (CLI Procedure) on page 2914
- Defining CoS Code-Point Aliases (J-Web Procedure) on page 2912
- Example: Configuring CoS on J-EX Series Switches on page 2883

Monitoring CoS Drop Profiles

Purpose Use the monitoring functionality to view data point information for each CoS random early detection (RED) drop profile on the J-EX8200 switch.

Action To monitor CoS RED drop profiles in the J-Web interface, select **Monitor > Class of Service > RED Drop Profiles**.

To monitor CoS RED drop profiles in the CLI, enter the following CLI command:

```
show class-of-service drop-profile
```

Meaning Table 384 on page 2941 summarizes the key output fields for CoS RED drop profiles.

Table 384: Summary of the Key Output Fields for CoS Red Drop Profiles

Field	Values	Additional Information
RED Drop Profile Name	Name of the RED drop profile. A drop profile consists of pairs of values between 0 and 100, one for queue buffer fill level and the other for drop probability, that determine the relationship between a buffer's fullness and the likelihood it will drop packets.	To display profile values, click the plus sign (+).
Graph RED Profile	Links to a graph of a RED curve that the system uses to determine the drop probability based on queue buffer fullness.	The x axis represents the queue buffer fill level, and the y axis represents the drop probability.

Table 384: Summary of the Key Output Fields for CoS Red Drop Profiles *(continued)*

Field	Values	Additional Information
Type	Type of a specific drop profile: • **interpolated**—The two coordinates (x and y) of the graph are interpolated to produce a smooth profile. • **segmented**—The two coordinates (x and y) of the graph are represented by line fragments to produce a segmented profile.	
Index	Internal index of this drop profile.	
Fill Level	Percentage fullness of a buffer queue. This value is the x coordinate of the RED drop profile graph.	
Drop Probability	Drop probability of a packet corresponding to a specific queue buffer fill level. This value is the y coordinate of the RED drop profile graph.	

Related Documentation

- Defining CoS Drop Profiles (J-Web Procedure) on page 2923
- Example: Configuring CoS on J-EX Series Switches on page 2883

CHAPTER 111

Configuration Statements for CoS

[edit class-of-service] Configuration Statement Hierarchy

```
class-of-service {
  classifiers {
    (dscp | ieee-802.1 | inet-precedence) classifier-name {
      import (classifier-name | default);
      forwarding-class class-name {
        loss-priority loss-priority {
          code-points [ aliases ] [ 6 bit-patterns ];
        }
      }
    }
  }
  code-point-aliases {
    (dscp | ieee-802.1 | inet-precedence) {
      alias-name bits;
    }
  }
  forwarding-classes {
    class class-name queue-num queue-number priority ( high | low );
  }
  interfaces {
    interface-name  {
      scheduler-map map-name;
      unit logical-unit-number {
        forwarding-class class-name;
        classifiers {
          (dscp | ieee-802.1 | inet-precedence) (classifier-name | default);
        }
      }
    }
  }
  multi-destination {
    family {
      ethernet {
        broadcast forwarding-class-name;
      }
      inet {
        classifiers {
```

```
                (dscp |inet-precedence) classifier-name;
            }
          }
        }
        scheduler-map  map-name;
      }
      rewrite-rules {
        (dscp | ieee-802.1 | inet-precedence) rewrite-name {
          import (rewrite-name | default);
          forwarding-class class-name {
            loss-priority loss-priority code-point (alias | bits);
          }
        }
      }
      scheduler-maps {
        map-name {
          forwarding-class class-name scheduler scheduler-name;
        }
      }
      schedulers {
        scheduler-name  {
          buffer-size (percent percentage | remainder);
          drop-profile-map loss-priority loss-priority protocol protocol drop-profile
            profile-name;
          priority priority;
          shaping-rate (rate | percent percentage);
          transmit-rate (rate | percent percentage | remainder);
        }
      }
    }
```

Related Documentation

- Example: Configuring CoS on J-EX Series Switches on page 2883
- Defining CoS Code-Point Aliases (CLI Procedure) on page 2914 or Defining CoS Code-Point Aliases (J-Web Procedure) on page 2912
- Defining CoS Classifiers (CLI Procedure) on page 2914 or Defining CoS Classifiers (J-Web Procedure) on page 2916
- Defining CoS Forwarding Classes (CLI Procedure) on page 2918 or Defining CoS Forwarding Classes (J-Web Procedure) on page 2918
- Configuring CoS Tail Drop Profiles (CLI Procedure) on page 2925
- Defining CoS Schedulers (CLI Procedure) on page 2920 or Defining CoS Schedulers (J-Web Procedure) on page 2920
- Defining CoS Rewrite Rules (CLI Procedure) on page 2925 or Defining CoS Rewrite Rules (J-Web Procedure) on page 2926
- Assigning CoS Components to Interfaces (CLI Procedure) on page 2928 or Assigning CoS Components to Interfaces (J-Web Procedure) on page 2928

broadcast

Syntax broadcast *forwarding-class-name*;

Hierarchy Level [edit **class-of-service multi-destination family ethernet**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the forwarding class for the broadcast traffic belonging to the Ethernet family.

Options ***forwarding-class-name*** —Name of the forwarding class:

- **mcast-af**—Default forwarding class for assured forwarding of multicast traffic.
- **mcast-be**—Default best-effort forwarding class for multicast traffic.
- **mcast-ef**—Default forwarding class for expedited forwarding of multicast traffic.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Understanding CoS Schedulers on page 2873
- Understanding CoS Forwarding Classes on page 2870
- Understanding CoS Classifiers on page 2867

buffer-size

Syntax buffer-size (exact | percent *percentage* | remainder);

Hierarchy Level [edit **class-of-service schedulers** *scheduler-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify buffer size.

Default If you do not include this statement, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 0, 0, 0, 0, and 5 percent.

Options

exact—Enforce the exact buffer size. When this option is configured, sharing is disabled on the queue, restricting the usage to guaranteed buffers only.

percent*percentage*—Buffer size as a percentage of total buffer.

remainder—Remaining buffer available.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring CoS on J-EX Series Switches on page 2883
- Defining CoS Schedulers (CLI Procedure) on page 2920 or Defining CoS Schedulers (J-Web Procedure) on page 2920
- Understanding CoS Schedulers on page 2873

class

Syntax class *class-name* queue-num *queue-number* priority (high | low);

Hierarchy Level [edit **class-of-service forwarding-classes**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure up to 16 forwarding classes with multiple forwarding classes mapped to single queues. If you want to configure up to eight forwarding classes with one-to-one mapping to output queues, use the **queue** statement instead of the **class** statement at the **[edit class-of-service forwarding-classes]** hierarchy level.

On J-EX8200 switches, you can assign a fabric priority to a forwarding class. The fabric priority determines whether packets belonging to the forwarding class are sent to the high priority ingress queue or the low priority ingress queue. By default, packets are sent to the low priority ingress queue. The primary use of this option is to prevent high priority input traffic from being dropped due to congestion on the port groups of oversubscribed line cards.

Options ***class-name***—Name of forwarding class.

queue-num ***queue-number***—Output queue number.
Range: 0 through 15.

priority (high | low)—(Optional) (J-EX8200 switches only) Fabric priority.
Values: **high** or **low**
Default: **low**

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring CoS on J-EX Series Switches on page 2883
- Defining CoS Forwarding Classes (CLI Procedure) on page 2918 or Defining CoS Forwarding Classes (J-Web Procedure) on page 2918

class-of-service

Syntax

```
class-of-service {
  classifiers {
    (dscp | ieee-802.1 | inet-precedence) classifier-name {
      import (classifier-name | default);
      forwarding-class class-name {
        loss-priority level {
          code-points [ aliases ] [ 6 bit-patterns ];
        }
      }
    }
  }
  code-point-aliases {
    (dscp | ieee-802.1 | inet-precedence) {
      alias-name  bits;
    }
  }
  forwarding-classes {
    class class-name queue-num queue-number priority ( high | low );
  }
  interfaces {
    interface-name  {
      scheduler-map  map-name;
      unit logical-unit-number {
        forwarding-class class-name;
        classifiers {
          (dscp | ieee-802.1 | inet-precedence) (classifier-name | default);
        }
      }
    }
  }
  multi-destination {
    family {
      ethernet {
        broadcast forwarding-class-name;
      }
      inet {
        classifiers {
          (dscp |inet-precedence) classifier-name;
        }
      }
    }
    scheduler-map  map-name;
  }
  rewrite-rules {
    (dscp | ieee-802.1 | inet-precedence) rewrite-name {
      import (rewrite-name | default);
      forwarding-class class-name {
        loss-priority priority code-point (alias | bits);
      }
    }
  }
  scheduler-maps {
```

```
    map-name {
      forwarding-class class-name scheduler scheduler-name;
    }
  }
  schedulers {
    scheduler-name  {
      buffer-size (percent percentage | remainder);
      drop-profile-map loss-priority loss-priority protocol protocol drop-profile profile-name;
      priority priority;
      shaping-rate (rate | percent percentage);
      transmit-rate (rate | percent percentage | remainder);
    }
  }
}
```

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure class-of-service (CoS) parameters on J-EX Series switches.

The remaining statements are explained separately.

Default If you do not configure any CoS features, the default CoS settings are used.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring CoS on J-EX Series Switches on page 2883
- Defining CoS Code-Point Aliases (CLI Procedure) on page 2914 or Defining CoS Code-Point Aliases (J-Web Procedure) on page 2912
- Defining CoS Classifiers (CLI Procedure) on page 2914 or Defining CoS Classifiers (J-Web Procedure) on page 2916
- Defining CoS Forwarding Classes (CLI Procedure) on page 2918 or Defining CoS Forwarding Classes (J-Web Procedure) on page 2918
- Configuring CoS Tail Drop Profiles (CLI Procedure) on page 2925
- Defining CoS Schedulers (CLI Procedure) on page 2920 or Defining CoS Schedulers (J-Web Procedure) on page 2920
- Defining CoS Rewrite Rules (CLI Procedure) on page 2925 or Defining CoS Rewrite Rules (J-Web Procedure) on page 2926
- Assigning CoS Components to Interfaces (CLI Procedure) on page 2928 or Assigning CoS Components to Interfaces (J-Web Procedure) on page 2928

classifiers

Syntax

```
classifiers {
  (dscp | ieee-802.1 | inet-precedence | exp) classifier-name {
    import (classifier-name | default);
    forwarding-class class-name {
      loss-priority level {
        code-points [ aliases ] [ 6–bit-patterns ];
      }
    }
  }
}
```

Hierarchy Level

[edit class-of-service],
[edit class-of-service interfaces *interface-name* unit *logical-unit-number*]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Apply a CoS aggregate behavior classifier to a logical interface. You can apply a default classifier or a custom classifier.

The statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring CoS on J-EX Series Switches on page 2883
- Example: Combining CoS with MPLS on J-EX Series Switches on page 2898
- Defining CoS Classifiers (CLI Procedure) on page 2914 or Defining CoS Classifiers (J-Web Procedure) on page 2916
- Assigning CoS Components to Interfaces (CLI Procedure) on page 2928 or Assigning CoS Components to Interfaces (J-Web Procedure) on page 2928
- Understanding CoS Classifiers on page 2867

code-point-aliases

Syntax

```
code-point-aliases {
  (dscp | ieee-802.1 | inet-precedence) |{
    alias-name  bits;
  }
}
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define an alias for a CoS marker.

The statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring CoS on J-EX Series Switches on page 2883
- Defining CoS Code-Point Aliases (CLI Procedure) on page 2914 or Defining CoS Code-Point Aliases (J-Web Procedure) on page 2912
- Understanding CoS Code-Point Aliases on page 2864

code-points

Syntax code-points [*aliases*] [*6 bit-patterns*];

Hierarchy Level [edit class-of-service classifiers (dscp | ieee-802.1 | inet-precedence) forwarding-class *class-name* loss-priority *level*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify one or more DSCP code-point aliases or bit sets for association with a forwarding class.

Options ***aliases*** —Name of the DSCP alias.

6 bit-patterns —Value of the code-point bits, in decimal form.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring CoS on J-EX Series Switches on page 2883
- Defining CoS Classifiers (CLI Procedure) on page 2914 or Defining CoS Classifiers (J-Web Procedure) on page 2916
- Understanding CoS Classifiers on page 2867

drop-profile-map

Syntax drop-profile-map **loss-priority** *loss-priority* **protocol** *protocol* drop-profile *profile-name*;

Hierarchy Level [edit **class-of-service schedulers** *scheduler-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define the loss priority value for the specified drop profile.

Options **drop-profile** ***profile-name*** —Name of the drop profile.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring CoS on J-EX Series Switches on page 2883
- Defining CoS Schedulers (CLI Procedure) on page 2920 or Defining CoS Schedulers (J-Web Procedure) on page 2920
- Understanding CoS Schedulers on page 2873

dscp

Syntax

```
dscp classifier-name {
  import (classifier-name | default);
  forwarding-class class-name {
    loss-priority level {
      code-points [ aliases ] [ 6–bit-patterns ];
    }
  }
}
```

Hierarchy Level

[edit class-of-service classifiers],
[edit class-of-service code-point-aliases],
[editclass-of-service interfaces *interface-name* unit *logical-unit-number* classifiers],
[edit class-of-service rewrite-rules]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Define the Differentiated Services code point (DSCP) mapping that is applied to the packets.

Options

classifier-name—Name of the classifier.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring CoS on J-EX Series Switches on page 2883
- Defining CoS Code-Point Aliases (CLI Procedure) on page 2914 or Defining CoS Code-Point Aliases (J-Web Procedure) on page 2912
- Defining CoS Classifiers (CLI Procedure) on page 2914 or Defining CoS Classifiers (J-Web Procedure) on page 2916
- Defining CoS Rewrite Rules (CLI Procedure) on page 2925 or Defining CoS Rewrite Rules (J-Web Procedure) on page 2926
- Assigning CoS Components to Interfaces (CLI Procedure) on page 2928 or Assigning CoS Components to Interfaces (J-Web Procedure) on page 2928
- Understanding CoS Classifiers on page 2867

dscp-ipv6

Syntax

```
dscp-ipv6 classifier-name {
  import (classifier-name | default);
  forwarding-classclass-name {
    loss-priority level {
      code-points [ aliases ] [ 6–bit-patterns ];
    }
  }
}
```

Hierarchy Level

```
[edit class-of-serviceclassifiers],
  [edit class-of-servicecode-point-aliases],
  [editclass-of-serviceinterfacesinterface-nameunitlogical-unit-number classifiers]
  [editclass-of-serviceinterfacesinterface-nameunitlogical-unit-number rewrite-rules]
  [edit class-of-servicerewrite-rules]
```

Release Information Statement introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Define the Differentiated Services code point (DSCP) mapping that is applied to the IPv6 packets.

Options *classifier-name*—Name of the classifier.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring CoS on J-EX Series Switches on page 2883
- Defining CoS Code-Point Aliases (CLI Procedure) on page 2914 or Defining CoS Code-Point Aliases (J-Web Procedure) on page 2912
- Defining CoS Classifiers (CLI Procedure) on page 2914 or Defining CoS Classifiers (J-Web Procedure) on page 2916
- Defining CoS Rewrite Rules (CLI Procedure) on page 2925 or Defining CoS Rewrite Rules (J-Web Procedure) on page 2926
- Assigning CoS Components to Interfaces (CLI Procedure) on page 2928 or Assigning CoS Components to Interfaces (J-Web Procedure) on page 2928
- Understanding CoS Classifiers on page 2867

ethernet

Syntax

```
ethernet {
  broadcast forwarding-class-name;
}
```

Hierarchy Level [edit **class-of-service multi-destination family**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the Ethernet broadcast traffic family.

The remaining statement is explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Understanding CoS Schedulers on page 2873
- Understanding CoS Forwarding Classes on page 2870
- Understanding CoS Classifiers on page 2867

exp

Syntax

```
exp classifier-name {
  import (classifier-name | default);
  forwarding-class class-name {
    loss-priority level {
      code-points [ aliases ] [ 3–bit-patterns ];
    }
  }
}
```

Hierarchy Level [edit class-of-service classifiers]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define the experimental bits (EXP) code point mapping that is applied to the MPLS packets.

J-EX Series switches support only one EXP code mapping on the switch (either default or custom). It is applied globally and implicitly to all the MPLS-enabled interfaces on the switch. You cannot bind it to an individual interface and you cannot disable it.

Options *classifier-name*—Name of the classifier.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Understanding Using CoS with MPLS Networks on J-EX Series Switches on page 2880
- Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect (CLI Procedure) on page 3111
- Configuring MPLS on Provider Edge Switches Using IP Over MPLS (CLI Procedure) on page 3107
- Configuring CoS on Provider Switches of an MPLS Network (CLI Procedure) on page 3106

family

Syntax

```
family {
  ethernet {
    broadcast forwarding-class-name;
  }
  inet {
    classifiers{
      (dscp | ieee-802.1 | inet-precedence) classifier-name;
    }
  }
}
```

Hierarchy Level [edit class-of-service multi-destination]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the multidestination traffic family.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Understanding CoS Schedulers on page 2873
- Understanding CoS Forwarding Classes on page 2870
- Understanding CoS Classifiers on page 2867

forwarding-class

Syntax

```
forwarding-class class-name {
  loss-priority level {
    code-points [ aliases ] [ 6–bit-patterns ];
  }
}
```

Hierarchy Level

[edit **class-of-service classifiers** (**dscp** | **ieee-802.1** | **inet-precedence**) *classifier-name*],
[edit**class-of-service interfaces** *interface-name* **unit** *logical-unit-number*],
[edit **class-of-service rewrite-rules**] (**dscp** | **ieee-802.1** | **inet-precedence**) *rewrite-name*],
[edit **class-of-service scheduler-maps** *map-name*]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Define forwarding class name and option values.

Options

class-name —Name of the forwarding class.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring CoS on J-EX Series Switches on page 2883
- Defining CoS Forwarding Classes (CLI Procedure) on page 2918 or Defining CoS Forwarding Classes (J-Web Procedure) on page 2918
- Understanding CoS Forwarding Classes on page 2870

forwarding-classes

Syntax

```
forwarding-classes {
  class class-name  queue-num queue-number;
}
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Associate the forwarding class with a queue name and number.

The statement is explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring CoS on J-EX-series Switches on page 2883
- Defining CoS Forwarding Classes (CLI Procedure) on page 2918 or Defining CoS Forwarding Classes (J-Web Procedure) on page 2918
- Understanding CoS Forwarding Classes on page 2870

ieee-802.1

Syntax

```
ieee-802.1 classifier-name {
  import (classifier-name | default);
  forwarding-class class-name {
    loss-priority level {
      code-points [ aliases ] [ 6 bit-patterns ];
    }
  }
}
```

Hierarchy Level

```
[edit class-of-service classifiers],
  [edit class-of-service code-point-aliases],
  [editclass-of-service interfaces interface-name unit logical-unit-number classifiers],
  [edit class-of-service rewrite-rules]
```

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply an IEEE-802.1 rewrite rule.

Options ***classifier-name*** —Name of the classifier.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring CoS on J-EX Series Switches on page 2883
- Defining CoS Classifiers (CLI Procedure) on page 2914 or Defining CoS Classifiers (J-Web Procedure) on page 2916
- Defining CoS Code-Point Aliases (CLI Procedure) on page 2914 or Defining CoS Code-Point Aliases (J-Web Procedure) on page 2912
- Defining CoS Rewrite Rules (CLI Procedure) on page 2925 or Defining CoS Rewrite Rules (J-Web Procedure) on page 2926
- Understanding CoS Classifiers on page 2867
- Understanding CoS Rewrite Rules on page 2876

import

Syntax import (*classifier-name* | default);

Hierarchy Level [edit **class-of-service classifiers** (**dscp** | **ieee-802.1** | **inet-precedence**) *classifier-name*],
[edit **class-of-service rewrite-rules** (**dscp** | **ieee-802.1** | **inet-precedence**) *rewrite-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify a default or previously defined classifier.

Options ***classifier-name*** —Name of the classifier mapping configured at the **[edit class-of-service classifiers]** hierarchy level.

default—Default classifier mapping.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring CoS on J-EX Series Switches on page 2883
- Defining CoS Classifiers (CLI Procedure) on page 2914 or Defining CoS Classifiers (J-Web Procedure) on page 2916
- Defining CoS Rewrite Rules (CLI Procedure) on page 2925 or Defining CoS Rewrite Rules (J-Web Procedure) on page 2926
- Understanding CoS Classifiers on page 2867
- Understanding CoS Rewrite Rules on page 2876

inet

Syntax

```
inet {
  classifiers {
    (dscp | ieee-802.1 | inet-precedence) classifier-name ;
  }
}
```

Hierarchy Level [edit class-of-service multi-destination family]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the IP multicast family.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Understanding CoS Schedulers on page 2873
- Understanding CoS Forwarding Classes on page 2870
- Understanding CoS Classifiers on page 2867

inet-precedence

Syntax

```
inet-precedence classifier-name {
  import (classifier-name | default);
  forwarding-class class-name {
    loss-priority level {
      code-points [ aliases ] [ 6–bit-patterns ];
    }
  }
}
```

Hierarchy Level

[edit **class-of-service classifiers**],
[edit **class-of-service code-point-aliases**],
[edit**class-of-service interfaces** *interface-name* **unit** *logical-unit-number* **classifiers**],
[edit **class-of-service rewrite-rules**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply an IPv4 precedence rewrite rule.

Options *classifier-name*—Name of the classifier.

The remaining statements are explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring CoS on J-EX Series Switches on page 2883
- Defining CoS Classifiers (CLI Procedure) on page 2914 or Defining CoS Classifiers (J-Web Procedure) on page 2916
- Defining CoS Code-Point Aliases (CLI Procedure) on page 2914 or Defining CoS Code-Point Aliases (J-Web Procedure) on page 2912
- Defining CoS Rewrite Rules (CLI Procedure) on page 2925 or Defining CoS Rewrite Rules (J-Web Procedure) on page 2926
- Understanding CoS Classifiers on page 2867
- Understanding CoS Rewrite Rules on page 2876

interfaces

Syntax

```
interfaces {
  interface-name  {
    scheduler-map map-name;
    unit logical-unit-number {
      forwarding-class class-name;
      classifiers {
        (dscp | ieee-802.1 | inet-precedence) (classifier-name | default);
      }
    }
  }
}
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure interface-specific CoS properties for incoming packets.

Options *interface-name* —Name of the interface.

The statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring CoS on J-EX Series Switches on page 2883
- Defining CoS Classifiers (CLI Procedure) on page 2914 or Defining CoS Classifiers (J-Web Procedure) on page 2916
- Defining CoS Forwarding Classes (CLI Procedure) on page 2918 or Defining CoS Forwarding Classes (J-Web Procedure) on page 2918
- Defining CoS Schedulers (CLI Procedure) on page 2920 or Defining CoS Schedulers (J-Web Procedure) on page 2920
- J-EX Series Switches Interfaces Overview on page 863

loss-priority

Syntax

```
loss-priority level {
  code-points [ aliases ] [ 6–bit-patterns | 3–bit-patterns ];
}
```

Hierarchy Level

[edit **class-of-service classifiers** (**dscp** | **ieee-802.1** | **inet-precedence** |exp) *classifier-name* **forwarding-class** *class-name*],
[edit **class-of-service rewrite-rules** (**dscp** | **ieee-802.1** | **inet-precedence** | exp) *rewrite-name* **forwarding-class** *class-name*]

Release Information

Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Specify packet loss priority value for a specific set of code-point aliases and bit patterns.

Options

level —Can be one of the following:

- **high**—Packet has high loss priority.
- **low**—Packet has low loss priority.

The remaining statement is explained separately.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring CoS on J-EX Series Switches on page 2883
- Defining CoS Classifiers (CLI Procedure) on page 2914 or Defining CoS Classifiers (J-Web Procedure) on page 2916
- Defining CoS Rewrite Rules (CLI Procedure) on page 2925 or Defining CoS Rewrite Rules (J-Web Procedure) on page 2926
- Understanding CoS Classifiers on page 2867
- Understanding CoS Rewrite Rules on page 2876

multi-destination

Syntax

```
multi-destination {
  family {
    ethernet {
      broadcast forwarding-class-name;
    }
    inet {
      classifiers {
        (dscp | ieee-802.1 | inet-precedence) classifier-name;
      }
    }
  }
  scheduler-map map-name;
}
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define the CoS configuration for multidestination traffic.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Understanding CoS Schedulers on page 2873
- Understanding CoS Forwarding Classes on page 2870
- Understanding CoS Classifiers on page 2867

policing

Syntax policing (filter *filter-name* | no-automatic-policing);

Hierarchy Level [edit protocols mpls label-switched-path *lsp-name*]
[edit interfaces *interface-id* unit *number-of-logical-unit* family inet address *ip-address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply a rate-limiting policer as the specified policing filter:

- To the LSP for MPLS over CCC.
- To the customer-edge interface for IP over MPLS.

Options **filter *filter-name***—Specify the name of the policing filter.

no-automatic-policing—Disable automatic policing on this LSP.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **policer on page 2826**
- Configuring Policers to Control Traffic Rates (CLI Procedure) on page 2788
- Configuring CoS on MPLS Provider Edge Switch Using Circuit Cross-Connect (CLI Procedure) on page 2932
- Configuring CoS on MPLS Provider Edge Switch Using IP Over MPLS (CLI Procedure) on page 2931

priority

Syntax priority *priority*;

Hierarchy Level [edit **class-of-service schedulers** *scheduler-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify packet-scheduling priority value.

Options ***priority*** —It can be one of the following:

- **low**—Scheduler has low priority.
- **strict-high**—Scheduler has strictly high priority.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring CoS on J-EX Series Switches on page 2883
- Defining CoS Schedulers (CLI Procedure) on page 2920 or Defining CoS Schedulers (J-Web Procedure) on page 2920
- Understanding CoS Schedulers on page 2873

protocol

Syntax protocol *protocol* drop-profile *profile-name*;

Hierarchy Level [edit **class-of-service schedulers** *scheduler-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the protocol type for the specified drop profile.

Options **drop-profile** ***profile-name*** —Name of the drop profile.

protocol —Type of protocol. It can be:

- **any**—Accept any protocol type.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring CoS on J-EX Series Switches on page 2883
- Configuring CoS Tail Drop Profiles (CLI Procedure) on page 2925
- Understanding CoS Tail Drop Profiles on page 2872

rewrite-rules

Syntax

```
rewrite-rules {
  (dscp | exp |ieee-802.1 | inet-precedence ) rewrite-name {
    import ( default | rewrite-name);
    forwarding-class class-name {
      loss-priority level code-point (alias | bits);
    }
  }
}
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify a rewrite-rules mapping for the traffic that passes through all queues on the interface.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Combining CoS with MPLS on J-EX Series Switches on page 2883
- Defining CoS Rewrite Rules (CLI Procedure) on page 2925 or Defining CoS Rewrite Rules (J-Web Procedure) on page 2926
- Understanding CoS Rewrite Rules on page 2876
- Understanding Using CoS with MPLS Networks on J-EX Series Switches on page 2880

scheduler-map

Syntax scheduler-map *map-name*;

Hierarchy Level [edit **class-of-service interfaces**],
[edit **class-of-service multi-destination**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Associate a scheduler map name with an interface or with a multidestination traffic configuration.

Options ***map-name*** —Name of the scheduler map.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring CoS on J-EX Series Switches on page 2883
- Assigning CoS Components to Interfaces (CLI Procedure) on page 2928 or Assigning CoS Components to Interfaces (J-Web Procedure) on page 2928
- Understanding CoS Schedulers on page 2873
- Understanding CoS Classifiers on page 2867

scheduler-maps

Syntax

```
scheduler-maps {
  map-name {
    forwarding-class class-name scheduler scheduler-name;
  }
}
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify a scheduler map name and associate it with the scheduler configuration and forwarding class.

Options *map-name* —Name of the scheduler map.

The remaining statement is explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring CoS on J-EX Series Switches on page 2883
- Defining CoS Forwarding Classes (CLI Procedure) on page 2918 or Defining CoS Forwarding Classes (J-Web Procedure) on page 2918
- Understanding CoS Schedulers on page 2873
- Understanding CoS Forwarding Classes on page 2870

schedulers

Syntax

```
schedulers {
  scheduler-name  {
    buffer-size (percent percentage | remainder);
    drop-profile-map loss-priority loss-priority protocol protocol drop-profile profile-name;
    priority priority;
    shaping-rate (rate | percent percentage);
    transmit-rate (rate | percent percentage | remainder);
  }
}
```

Hierarchy Level [edit class-of-service]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify scheduler name and parameter values.

Options ***scheduler-name*** —Name of the scheduler.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring CoS on J-EX Series Switches on page 2883
- Defining CoS Schedulers (CLI Procedure) on page 2920 or Defining CoS Schedulers (J-Web Procedure) on page 2920
- Understanding CoS Schedulers on page 2873

shaping-rate

Syntax shaping-rate (percent *percentage* | rate);

Hierarchy Level [edit **class-of-service schedulers** *scheduler-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure shaping rate to throttle the rate at which queues transmit packets.

We recommend that you configure the shaping rate as an absolute maximum usage and not as additional usage beyond the configured transmit rate.

Default If you do not include this statement, the default shaping rate is 100 percent, which is the same as no shaping at all.

Options **percent***percentage* —Shaping rate as a percentage of the available interface bandwidth.

Range: 0 through 100 percent

rate—Peak rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000).

Range: 3200 through 32,000,000,000 bps

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring CoS on J-EX Series Switches on page 2883
- Understanding Junos OS CoS Components for J-EX Series Switches on page 2862

shared-buffer

Syntax shared-buffer percent *percentage*

Hierarchy Level [edit **class-of-service**],

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the buffer allocation for the shared buffer pool.

Options **percent *percentage***—Size of the shared buffer as a percentage of the buffer allocated to the shared buffer pool.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring CoS on J-EX Series Switches on page 2883
- Understanding Junos OS CoS Components for J-EX Series Switches on page 2862

transmit-rate

Syntax transmit-rate (*rate* | percent *percentage* | remainder);

Hierarchy Level [edit **class-of-service schedulers** *scheduler-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the transmit rate or percentage for a scheduler.

Default If you do not include this statement, the default scheduler transmission rate and buffer size percentages for queues 0 through 7 are 95, 0, 0, 0, 0, 0, 0, and 5 percent.

Options

rate —Transmission rate, in bps. You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000).

Range: 3200 through 160,000,000,000 bps

percent *percentage* —Percentage of transmission capacity. A percentage of zero drops all packets in the queue.

Range: 0 through 100 percent

remainder—Remaining rate available

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring CoS on J-EX Series Switches on page 2883
- Defining CoS Schedulers (CLI Procedure) on page 2920 or Defining CoS Schedulers (J-Web Procedure) on page 2920
- Understanding CoS Schedulers on page 2873

unit

Syntax

```
unit logical-unit-number {
  forwarding-class class-name;
  classifiers {
    (dscp | ieee-802.1 | inet-precedence) (classifier-name | default);
  }
}
```

Hierarchy Level [edit class-of-service interfaces *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.

Options ***logical-unit-number*** —Number of the logical unit.

Range: 0 through 16,385

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring CoS on J-EX Series Switches on page 2883
- Assigning CoS Components to Interfaces (CLI Procedure) on page 2928 or Assigning CoS Components to Interfaces (J-Web Procedure) on page 2928

CHAPTER 112

Operational Mode Commands for CoS

show class-of-service

Syntax **show class-of-service**

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display the class-of-service (CoS) information.

Required Privilege Level view

Related Documentation

- Example: Configuring CoS on J-EX Series Switches on page 2883
- Monitoring CoS Value Aliases on page 2940
- Monitoring CoS Classifiers on page 2935
- Monitoring CoS Forwarding Classes on page 2936
- Monitoring CoS Scheduler Maps on page 2939
- Monitoring CoS Rewrite Rules on page 2938

List of Sample Output **show class-of- service on page 2979**
show class-of-service rewrite-rule on page 2982

Output Fields Table 385 on page 2978 lists the output fields for the **show class-of-service** command. Output fields are listed in the approximate order in which they appear.

Table 385: show class-of-service Output Fields

Field Name	Field Description	Level of Output
Forwarding class	The forwarding class configuration: • **Forwarding class**—Name of the forwarding class. • **ID**—Forwarding class ID. • **Queue**—Queue number. • **Fabric Priority**—(J-EX8200 switches only) Fabric priority: either **high** or **low**. The fabric priority determines which CoS ingress queues packets are sent to.	All levels
Code point type	The type of code-point alias: • **dscp**—Aliases for DiffServ code point (DSCP) values. • **ieee–802.1**—Aliases for IEEE 802.1p values. • **inet-precedence**—Aliases for IP precedence values. • **exp**—Aliases for experimental (EXP) values.	All levels
Alias	Names given to CoS values.	All levels
Bit pattern	Set of bits associated with an alias.	All levels
Classifier	Name of the classifier.	All levels

Table 385: show class-of-service Output Fields *(continued)*

Field Name	Field Description	Level of Output
Code point	Code-point values.	All levels
Loss priority	Loss priority assigned to specific CoS values and aliases of the classifier.	All levels
Rewrite rule	Name of the rewrite-rule.	All levels
Drop profile	Name of the drop profile.	All levels
Type	Type of drop profile. J-EX Series switches support only the **discrete** type of drop profile.	All levels
Fill level	Percentage of queue buffer fullness of high packets beyond which high packets are dropped.	All levels
Scheduler	Name of the scheduler.	All levels
Transmit rate	Transmission rate of the scheduler.	All levels
Buffer size	Delay buffer size in the queue.	All levels
Drop profiles	Drop profiles configured for the specified scheduler.	All levels
Protocol	Transport protocol corresponding to the drop profile.	All levels
Name	Name of the drop profile.	All levels
Queues supported	Number of queues that can be configured on the interface.	All levels
Queues in use	Number of queues currently configured.	All levels
Physical interface	Name of the physical interface.	All levels
Scheduler map	Name of the scheduler map.	All levels
Index	Internal index of a specific object.	All levels

show class-of- service

```
user@switch> show class-of-service
Forwarding class                          ID      Queue
  best-effort                              0       0
  expedited-forwarding                     1       5
  assured-forwarding                       2       1
  network-control                          3       7

Code point type: dscp
  Alias              Bit pattern
  af11               001010
  af12               001100
  ...                ...
```

```
Code point type: ieee-802.1
  Alias              Bit pattern
  af11               010
  ...                ...

Code point type: inet-precedence
  Alias              Bit pattern
  af11               001
  ...                ...

Classifier: dscp-default, Code point type: dscp, Index: 7
  Code point         Forwarding class                   Loss priority
  000000             best-effort                        low
  000001             best-effort                        low
  ...                ...                                ...

Classifier: ieee8021p-default, Code point type: ieee-802.1, Index: 11
  Code point         Forwarding class                   Loss priority
  000                best-effort                        low
  001                best-effort                        low
  010                best-effort                        low
  011                best-effort                        low
  100                best-effort                        low
  101                best-effort                        low
  110                network-control                    low
  111                network-control                    low

Classifier: ipprec-default, Code point type: inet-precedence, Index: 12
  Code point         Forwarding class                   Loss priority
  000                best-effort                        low
  001                best-effort                        low
  010                best-effort                        low
  011                best-effort                        low
  100                best-effort                        low
  101                best-effort                        low
  110                network-control                    low
  111                network-control                    low

Classifier: ieee8021p-untrust, Code point type: ieee-802.1, Index: 16
  Code point         Forwarding class                   Loss priority
  000                best-effort                        low
  001                best-effort                        low
  010                best-effort                        low
  011                best-effort                        low
  100                best-effort                        low
  101                best-effort                        low
  110                best-effort                        low
  111                best-effort                        low

Rewrite rule: dscp-default, Code point type: dscp, Index: 27
  Forwarding class                   Loss priority       Code point
  best-effort                        low                 000000
  best-effort                        high                000000
  expedited-forwarding               low                 101110
  expedited-forwarding               high                101110
  assured-forwarding                 low                 001010
  assured-forwarding                 high                001100
  network-control                    low                 110000
  network-control                    high                111000
```

```
Rewrite rule: ieee8021p-default, Code point type: ieee-802.1, Index: 30
  Forwarding class                  Loss priority       Code point
  best-effort                       low                 000
  best-effort                       high                001
  expedited-forwarding              low                 100
  expedited-forwarding              high                101
  assured-forwarding                low                 010
  assured-forwarding                high                011
  network-control                   low                 110
  network-control                   high                111

Rewrite rule: ipprec-default, Code point type: inet-precedence, Index: 31
  Forwarding class                  Loss priority       Code point
  best-effort                       low                 000
  best-effort                       high                000
  expedited-forwarding              low                 101
  expedited-forwarding              high                101
  assured-forwarding                low                 001
  assured-forwarding                high                001
  network-control                   low                 110
  network-control                   high                111

Drop profile:<default-drop-profile>, Type: discrete, Index: 1
  Fill level
         100

Scheduler map: <default>, Index: 2

  Scheduler: <default-be>, Forwarding class: best-effort, Index: 20
    Transmit rate: 95 percent, Rate Limit: none, Buffer size: 95 percent,
    Priority: low
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

  Scheduler: <default-nc>, Forwarding class: network-control, Index: 22
    Transmit rate: 5 percent, Rate Limit: none, Buffer size: 5 percent,
    Priority: low
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>

Physical interface: ge-0/0/0, Index: 129
Queues supported: 8, Queues in use: 4
  Scheduler map: <default>, Index: 2

Physical interface: ge-0/0/1, Index: 130
Queues supported: 8, Queues in use: 4
  Scheduler map: <default>, Index: 2

...             ...              ...

Fabric priority: low
  Scheduler: <default-fabric>, Index: 23
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>
```

```
Fabric priority: high
  Scheduler: <default-fabric>, Index: 23
    Drop profiles:
      Loss priority   Protocol    Index    Name
      High            non-TCP         1    <default-drop-profile>
      High            TCP             1    <default-drop-profile>
```

show class-of-service rewrite-rule

```
user@switch> show class-of-service rewrite-rule
Rewrite rule: dscp-default, Code point type: dscp, Index: 31
  Forwarding class                  Loss priority       Code point
  best-effort                       low                 000000
  best-effort                       high                000000
  expedited-forwarding              low                 101110
  expedited-forwarding              high                101110
  fw-class                          low                 001010
  fw-class                          high                001100
  network-control                   low                 110000
  network-control                   high                111000

Rewrite rule: exp-default, Code point type: exp, Index: 33
  Forwarding class                  Loss priority       Code point
  best-effort                       low                 000
  best-effort                       high                001
  expedited-forwarding              low                 010
  expedited-forwarding              high                011
  fw-class                          low                 100
  fw-class                          high                101
  network-control                   low                 110
  network-control                   high                111

Rewrite rule: ieee8021p-default, Code point type: ieee-802.1, Index: 34
  Forwarding class                  Loss priority       Code point
  best-effort                       low                 000
  best-effort                       high                001
  expedited-forwarding              low                 010
  expedited-forwarding              high                011
  fw-class                          low                 100
  fw-class                          high                101
  network-control                   low                 110
  network-control                   high                111

Rewrite rule: ipprec-default, Code point type: inet-precedence, Index: 35
  Forwarding class                  Loss priority       Code point
  best-effort                       low                 000
  best-effort                       high                000
  expedited-forwarding              low                 101
  expedited-forwarding              high                101
  fw-class                          low                 001
  fw-class                          high                001
  network-control                   low                 110
  network-control                   high                111
```

show class-of-service classifier

Syntax

show class-of-service classifier
<name *name*>
<type dscp | type dscp-ipv6 | type exp | type ieee-802.1 | type inet-precedence>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

For each class-of-service (CoS) classifier, display the mapping of code point value to forwarding class and loss priority.

Options

none—Display all classifiers.

name *name*—(Optional) Display named classifier.

type dscp—(Optional) Display all classifiers of the Differentiated Services code point (DSCP) type.

type dscp-ipv6—(Optional) Display all classifiers of the DSCP for IPv6 type.

type exp—(Optional) Display all classifiers of the MPLS experimental (EXP) type.

type ieee-802.1—(Optional) Display all classifiers of the ieee-802.1 type.

type inet-precedence—(Optional) Display all classifiers of the inet-precedence type.

Required Privilege Level

view

List of Sample Output

show class-of-service classifier type ieee-802.1 on page 2984

Output Fields

Table 386 on page 2983 describes the output fields for the **show class-of-service classifier** command. Output fields are listed in the approximate order in which they appear.

Table 386: show class-of-service classifier Output Fields

Field Name	Field Description
Classifier	Name of the classifier.
Code point type	Type of the classifier: **exp** (not on J-EX Series switches), **dscp**, **dscp-ipv6** (not on J-EX Series switches), **ieee-802.1**, or **inet-precedence**.
Index	Internal index of the classifier.
Code point	Code point value used for classification
Forwarding class	Classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router.

Table 386: show class-of-service classifier Output Fields *(continued)*

Field Name	Field Description
Loss priority	Loss priority value used for classification. For most platforms, the value is **high** or **low**. For some platforms, the value is **high**, **medium-high**, **medium-low**, or **low**.

show class-of-service classifier type ieee-802.1

```
user@host>  show class-of-service classifier type ieee-802.1
Classifier: ieee802.1-default, Code point type: ieee-802.1, Index: 3
Code Point         Forwarding Class                  Loss priority
  000              best-effort                         low
  001              best-effort                         high
  010              expedited-forwarding                low
  011              expedited-forwarding                high
  100              assured-forwarding                  low
  101              assured-forwarding                  medium-high
  110              network-control                     low
  111              network-control                     high

Classifier: users-ieee802.1, Code point type: ieee-802.1
  Code point         Forwarding class                  Loss priority
  100                expedited-forwarding              low
```

show class-of-service code-point-aliases

Syntax show class-of-service code-point-aliases
<dscp | dscp-ipv6 | exp | ieee-802.1 | inet-precedence>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display the mapping of class-of-service (CoS) code point aliases to corresponding bit patterns.

Options none—Display code point aliases of all code point types.

dscp—(Optional) Display Differentiated Services code point (DSCP) aliases.

dscp-ipv6—(Optional) Display IPv6 DSCP aliases.

exp—(Optional) Display MPLS EXP code point aliases.

ieee-802.1—(Optional) Display IEEE-802.1 code point aliases.

inet-precedence—(Optional) Display IPv4 precedence code point aliases.

Required Privilege Level view

List of Sample Output **show class-of-service code-point-aliases exp on page 2986**

Output Fields Table 387 on page 2985 describes the output fields for the **show class-of-service code-point-aliases** command. Output fields are listed in the approximate order in which they appear.

Table 387: show class-of-service code-point-aliases Output Fields

Field Name	Field Description
Code point type	Type of the code points displayed: **dscp**, **dscp-ipv6** (not on J-EX Series switches), **exp** (not on J-EX Series switches), **ieee-802.1**, or **inet-precedence**.
Alias	Alias for a bit pattern.
Bit pattern	Bit pattern for which the alias is displayed.

show class-of-service code-point-aliases exp

```
user@host> show class-of-service code-point-aliases exp
Code point type: exp
  Alias               Bit pattern
  af11                100
  af12                101
  be                  000
  be1                 001
  cs6                 110
  cs7                 111
  ef                  010
  ef1                 011
  nc1                 110
  nc2                 111
```

show class-of-service drop-profile

Syntax show class-of-service drop-profile
<profile-name *profile-name*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display data points for each class-of-service (CoS) random early detection (RED) drop profile.

Options none—Display all drop profiles.

profile-name *profile-name*—(Optional) Display the specified profile only.

Required Privilege Level view

List of Sample Output **show class-of-service drop-profile on page 2988**

Output Fields Table 388 on page 2987 describes the output fields for the **show class-of-service drop-profile** command. Output fields are listed in the approximate order in which they appear.

Table 388: show class-of-service drop-profile Output Fields

Field Name	Field Description
Drop profile	Name of a drop profile.
Type	Type of this drop profile: **discrete** or **interpolated**.
Index	Internal index of this drop profile.
Fill Level	Percentage fullness of a queue.
Drop probability	Drop probability at this fill level.

show class-of-service drop-profile

```
user@host> show class-of-service drop-profile
Drop profile: <default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
         100                 100
Drop profile: user-drop-profile, Type: interpolated, Index: 2989
  Fill level    Drop probability
           0                   0
           1                   1
           2                   2
           4                   4
           5                   5
           6                   6
           8                   8
          10                  10
          12                  15
          14                  20
          15                  23
        ... 64 entries total
          90                  96
          92                  96
          94                  97
          95                  98
          96                  98
          98                  99
          99                  99
         100                 100
```

show class-of-service forwarding-class

Syntax show class-of-service forwarding-class

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display information about forwarding classes, including the mapping of forwarding classes to queue numbers.

Required Privilege Level view

Related Documentation

- Example: Configuring CoS on J-EX Series Switches on page 2883
- Monitoring CoS Forwarding Classes on page 2936
- Defining CoS Forwarding Classes (CLI Procedure) on page 2918

List of Sample Output

show class-of-service forwarding-class on page 2989
show class-of-service forwarding-class (J-EX8200 Switch) on page 2989

Output Fields Table 389 on page 2989 describes the output fields for the **show class-of-service forwarding-class** command. Output fields are listed in the approximate order in which they appear.

Table 389: show class-of-service forwarding-class Output Fields

Field Name	Field Description
Forwarding class	Name of forwarding class.
ID	Forwarding class identifier.
Queue	CoS queue mapped to the forwarding class.
Policing priority	Not supported on J-EX Series switches and can be ignored.
Fabric priority	(J-EX8200 switches only) Fabric priority for the forwarding class, either **high** or **low**. Determines the priority of packets ingressing the switch fabric.

show class-of-service forwarding-class

```
user@switch> show class-of-service forwarding-class
Forwarding class                ID                      Queue  Policing priority
  best-effort                   0                           0           normal
  expedited-forwarding          1                           5           normal
  assured-forwarding            2                           1           normal
  network-control               3                           7           normal
```

show class-of-service forwarding-class (J-EX8200 Switch)

```
user@switch> show class-of-service forwarding-class
Forwarding class                  ID      Queue  Fabric priority
  best-effort                     0       0      low
  expedited-forwarding            1       5      low
  assured-forwarding              2       1      low
```

```
network-control                 3        7       low
mcast-be                        4        2       low
mcast-ef                        5        4       low
mcast-af                        6        6       low
```

show class-of-service interface

Syntax show class-of-service interface
<*interface-name*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display the logical and physical interface associations for the classifier, rewrite rules, and scheduler map objects.

Options none—Display class of service (CoS) associations for all physical and logical interfaces.

interface-name—(Optional) Display CoS associations for the specified interface.

Required Privilege Level view

List of Sample Output
show class-of-service interface (Physical) on page 2992
show class-of-service interface (Logical) on page 2992
show class-of-service interface (Gigabit Ethernet) on page 2993

Output Fields Table 390 on page 2991 describes the output fields for the **show class-of-service interface** command. Output fields are listed in the approximate order in which they appear.

Table 390: show class-of-service interface Output Fields

Field Name	Field Description
Physical interface	Name of a physical interface.
Index	Index of this interface or the internal index of this object.
Dedicated Queues	Status of dedicated queues configured on an interface. Supported on Trio MPC/MIC interfaces on MX Series routers only.
Queues supported	Number of queues you can configure on the interface.
Queues in use	Number of queues currently configured.
Total non-default queues created	Number of queues created in addition to the default queues. Supported on Trio MPC/MIC interfaces on MX Series routers.
Shaping rate	Maximum transmission rate on the physical interface. You can configure the shaping rate on the physical interface, or on the logical interface, but not both. Therefore, the **Shaping rate** field is displayed for the physical interface or the logical interface, but not both.
Scheduler map	Name of the output scheduler map associated with this interface.
Input shaping rate	For Gigabit Ethernet IQ2 PICs, maximum transmission rate on the input interface.
Input scheduler map	For Gigabit Ethernet IQ2 PICs, name of the input scheduler map associated with this interface.

Table 390: show class-of-service interface Output Fields *(continued)*

Field Name	Field Description
Chassis scheduler map	Name of the scheduler map associated with the packet forwarding component queues.
Rewrite	Name and type of the rewrite rules associated with this interface.
Classifier	Name and type of classifiers associated with this interface.
Forwaridng-class-map	Name of the forwarding map associated with this interface.
Logical interface	Name of a logical interface.
Shaping rate	Maximum transmission rate on the logical interface. You can configure the shaping rate on the physical interface, or on the logical interface, but not both. Therefore, the **Shaping rate** field is displayed for the physical interface or the logical interface, but not both.
Object	Category of an object: **Classifier**, **Fragmentation-map** (for LSQ interfaces only), **Scheduler-map**, **Rewrite**, or **Translation Table** (for IQE PICs only).
Name	Name of an object.
Type	Type of an object: **dscp**, **dscp-ipv6**, **exp**, **ieee-802.1**, **ip**, or **inet-precedence**.

show class-of-service interface (Physical)

```
user@host> show class-of-service interface so-0/2/3
Physical interface: so-0/2/3, Index: 135
Queues supported: 8, Queues in use: 4
Total non-default queues created: 4
  Scheduler map: <default>, Index: 2032638653

  Logical interface: fe-0/0/1.0, Index: 68, Dedicated Queues: no
    Shaping rate: 32000
    Object                   Name                   Type
Index
    Scheduler-map            <default>
  27
    Rewrite                  exp-default            exp
  21
    Classifier               exp-default            exp
   5
    Classifier               ipprec-compatibility   ip
   8
    Forwarding-class-map     exp-default            exp
   5
```

show class-of-service interface (Logical)

```
user@host> show class-of-service interface so-0/2/3.0
Logical interface: so-0/2/3.0, Index: 68, Dedicated Queues: no
    Shaping rate: 32000
    Object                   Name                   Type
Index
    Scheduler-map            <default>
  27
    Rewrite                  exp-default            exp
  21
```

```
  Classifier               exp-default            exp
5
  Classifier               ipprec-compatibility   ip
8
  Forwarding-class-map     exp-default            exp
5
```

show class-of-service interface (Gigabit Ethernet)

```
user@host> show class-of-service interface ge-6/2/0
Physical interface: ge-6/2/0, Index: 175
Queues supported: 4, Queues in use: 4
  Scheduler map: <default>, Index: 2
  Input scheduler map: <default>, Index: 3
  Chassis scheduler map: <default-chassis>, Index: 4
```

show pfe statistics traffic

Syntax show pfe statistics traffic

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display the packet forwarding engine traffic statistics.

Options none—Display statistics about all the traffic handled by the packet forwarding engine.

Required Privilege Level admin

List of Sample Output **show pfe statistics traffic on page 2995**

Output Fields Table 391 on page 2994 lists the output fields for the **show pfe statistics traffic** command. Output fields are listed in the approximate order in which they appear.

Table 391: show pfe statistics traffic Output Fields

Field Name	Field Description
Packet Forwarding Engine Traffic statistics	Information about Packet Forwarding Engine traffic: • **Input Packets**—Number and rate of input packets. • **Output Packets**—Number and rate of output packets.
Packet Forwarding Engine Local Traffic statistics	Information about Packet Forwarding Engine local traffic: • **Local packets input**—Number of local input packets. • **Local packets output**—Number of local output packets. • **Software input high drops**—Number of software input high-priority drops. • **Software input medium drops**—Number of software input medium-priority drops. • **Software input low drops**—Number of software input low-priority drops. • **Software output drops**—Number of software output drops. • **Hardware input drops**—Number of hardware input drops.

Table 391: show pfe statistics traffic Output Fields *(continued)*

Field Name	Field Description
Packet Forwarding Engine Local Protocol statistics	Information about the Packet Forwarding Engine Local Protocol: • **HDLC keepalives**—Number of HDLC keepalive packets. • **ATM OAM**—Number of Asynchronous Transfer Mode (ATM) Operation, Administration, and Maintenance (OAM) packets. • **Frame Relay LMI**—Number of Frame Relay Local Management Interface (LMI) packets. • **PPP LCP/NCP**—Number of Point-to-Point Protocol (PPP) Link Control Protocol (LCP) or Network Control Protocol (NCP) packets. • **OSPF hello**—Number of Open Shortest Path First (OSPF) hello packets. • **OSPF3 hello**—Number of Open Shortest Path First version 3 (OSPFv3) hello packets. • **RSVP hello**—Number of Reservation Setup Protocol (RSVP) hello packets. • **LDP hello**—Number of Label Distribution Protocol (LDP) hello packets. • **BFD**—Number of Bidirectional Forwarding Detection Protocol (BFD) hello packets. • **IS-IS IIH**—Number of Intermediate System-to-Intermediate System Hello (IIH) packets. • **LACP**—Number of Link Aggregation Control Protocol (LACP) packets. • **ARP**—Number of Address Resolution Protocol (ARP) packets. • **ETHER OAM**—Number of Ethernet Operations, Administration, and Management (OAM) packets. • **Unknown**—Number of unknown packets not matching any of the packet types listed above.
Packet Forwarding Engine Hardware Discard statistics	Information about Packet Forwarding Engine hardware discards: • **Timeout**—Number of packets discarded because of timeouts. • **Truncated key**—Number of packets discarded because of truncated keys. • **Bits to test**—Number of bits to test. • **Data error**—Number of packets discarded because of data errors. • **Stack underflow**—Number of packets discarded because of stack underflows. • **Stack overflow**—Number of packets discarded because of stack overflows. • **Normal discard**—Number of packets discarded because of discard routes. • **Extended discard**—Number of packets discarded because of illegal next hops. • **Invalid interface**—Number of packets discarded because of invalid incoming interfaces. • **Info cell drops**—Number of information cell drops. • **Fabric drops**—Number of fabric drops.

show pfe statistics traffic

```
user@host> show pfe statistics traffic
Packet Forwarding Engine traffic statistics:
    Input  packets:              102682                    5 pps
    Output packets:               58033                    4 pps
Packet Forwarding Engine local traffic statistics:
    Local packets input                :               44628
    Local packets output               :               46146
    Software input control plane drops  :                  0
    Software input high drops          :                  0
    Software input medium drops        :                  0
    Software input low drops           :                  0
    Software output drops              :                  0
    Hardware input drops               :                  0
Packet Forwarding Engine local protocol statistics:
    HDLC keepalives            :                   0
```

```
    ATM OAM                    :                    0
    Frame Relay LMI            :                    0
    PPP LCP/NCP                :                 5597
    OSPF hello                 :                 3195
    OSPF3 hello                :                    0
    RSVP hello                 :                    0
    LDP hello                  :                 7478
    BFD                        :                    0
    IS-IS IIH                  :                    0
    LACP                       :                    0
    ARP                        :                    0
    ETHER OAM                  :                    0
    Unknown                    :                    8
Packet Forwarding Engine hardware discard statistics:
    Timeout                    :                    0
    Truncated key              :                    0
    Bits to test               :                    0
    Data error                 :                    0
    Stack underflow            :                    0
    Stack overflow             :                    0
    Normal discard             :                    0
    Extended discard           :                    0
    Invalid interface          :                    0
    Info cell drops            :                    0
    Fabric drops               :                    0
Packet Forwarding Engine Input IPv4 Header Checksum Error and Output MTU Error
statistics:
    Input Checksum             :                    0
    Output MTU                 :                    0
```

show pfe statistics traffic cpu

Syntax show pfe statistics traffic cpu <fpc *fpc-slot*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description (On J-EX8200 switches only) Display count of multidestination packets ingressing from the physical interface to the CPU.

NOTE: Multidestination packets include unknown unicast, broadcast, and multicast packets.

Options none—Displays the count of packets ingressing from all the physical interfaces (line cards) to the CPU.

fpc *fpc-slot*—(Optional) Displays the count of packets ingressing from the physical interface, referred to by the slot number, to the CPU.

On a J-EX8200 switch, the FPC slot number is the slot number for the line card. Possible values are **0** through **7** on the J-EX8208 switch and **0** through **15** on the J-EX8216 switch.

Required Privilege Level view

Related Documentation

- **show pfe statistics traffic multicast on page 3003**
- **show pfe statistics traffic egress-queues on page 3001**
- **show interfaces queue on page 1016**
- Monitoring Interface Status and Traffic on page 931
- Understanding Junos OS CoS Components for J-EX Series Switches on page 2862

List of Sample Output **show pfe statistics traffic cpu (J-EX8208 Switch) on page 2998**

Output Fields Table 392 on page 2997 lists the output fields for the **show pfe statistics traffic cpu** command. Output fields are listed in the approximate order in which they appear.

Table 392: show pfe statistics traffic cpu Output Fields

Field Name	Field Description
Queue	CoS queue number.
Forwarding classes	Forwarding class name.
Queued Packets	Number of packets queued to this queue.
Queued Bytes	Number of bytes queued to this queue.

Table 392: show pfe statistics traffic cpu Output Fields *(continued)*

Field Name	Field Description
Packets	Number of packets transmitted by this queue.
Bytes	Number of bytes transmitted by this queue.
Tail-dropped packets	Count of packets dropped at the tail end of the queue because of lack of buffer space.
RED-dropped packets	Number of packets dropped because of Random Early Discard (RED): • **Low**—Number of low-loss priority packets dropped because of RED. • **High**—Number of high-loss priority packets dropped because of RED.
RED-dropped bytes	Number of bytes dropped because of Random Early Discard (RED): • **Low**—Number of low-loss priority bytes dropped because of RED. • **High**—Number of high-loss priority bytes dropped because of RED.

show pfe statistics traffic cpu (J-EX8208 Switch)

```
user@switch> show pfe statistics traffic cpu

Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              : Not Available
    Bytes                : Not Available
    Packets              :                       0                       0 pps
    Bytes                :                       0                       0 bps
    Tail-dropped packets :                       0
    RED-dropped bytes    :                       0                       0 bps
     Low                 :                       0                       0 bps
     High                :                       0                       0 bps
    RED-dropped packets  :                       0                       0 pps
     Low                 :                       0                       0 pps
     High                :                       0                       0 pps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              : Not Available
    Bytes                : Not Available
    Packets              :                       0                       0 pps
    Bytes                :                       0                       0 bps
    Tail-dropped packets :                       0
    RED-dropped bytes    :                       0                       0 bps
     Low                 :                       0                       0 bps
     High                :                       0                       0 bps
    RED-dropped packets  :                       0                       0 pps
     Low                 :                       0                       0 pps
     High                :                       0                       0 pps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              : Not Available
    Bytes                : Not Available
    Packets              :                       0                       0 pps
    Bytes                :                       0                       0 bps
    Tail-dropped packets :                       0
    RED-dropped bytes    :                       0                       0 bps
     Low                 :                       0                       0 bps
     High                :                       0                       0 bps
```

```
    RED-dropped packets  :                     0                    0 pps
     Low                 :                     0                    0 pps
     High                :                     0                    0 pps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets              : Not Available
    Bytes                : Not Available
    Packets              :                     0                    0 pps
    Bytes                :                     0                    0 bps
    Tail-dropped packets :                     0
    RED-dropped bytes    :                     0                    0 bps
     Low                 :                     0                    0 bps
     High                :                     0                    0 bps
    RED-dropped packets  :                     0                    0 pps
     Low                 :                     0                    0 pps
     High                :                     0                    0 pps
Queue: 4
    Packets              : Not Available
    Bytes                : Not Available
    Packets              :                     0                    0 pps
    Bytes                :                     0                    0 bps
    Tail-dropped packets :                     0
    RED-dropped bytes    :                     0                    0 bps
     Low                 :                     0                    0 bps
     High                :                     0                    0 bps
    RED-dropped packets  :                     0                    0 pps
     Low                 :                     0                    0 pps
     High                :                     0                    0 pps
Queue: 5
    Packets              : Not Available
    Bytes                : Not Available
    Packets              :                     0                    0 pps
    Bytes                :                     0                    0 bps
    Tail-dropped packets :                     0
    RED-dropped bytes    :                     0                    0 bps
     Low                 :                     0                    0 bps
     High                :                     0                    0 bps
    RED-dropped packets  :                     0                    0 pps
     Low                 :                     0                    0 pps
     High                :                     0                    0 pps
Queue: 6
    Packets              : Not Available
    Bytes                : Not Available
    Packets              :                     0                    0 pps
    Bytes                :                     0                    0 bps
    Tail-dropped packets :                     0
    RED-dropped bytes    :                     0                    0 bps
     Low                 :                     0                    0 bps
     High                :                     0                    0 bps
    RED-dropped packets  :                     0                    0 pps
     Low                 :                     0                    0 pps
     High                :                     0                    0 pps
Queue: 7
    Packets              : Not Available
    Bytes                : Not Available
    Packets              :                     0                    0 pps
    Bytes                :                     0                    0 bps
    Tail-dropped packets :                     0
    RED-dropped bytes    :                     0                    0 bps
     Low                 :                     0                    0 bps
     High                :                     0                    0 bps
```

```
RED-dropped packets  :                0                0 pps
 Low                 :                0                0 pps
 High                :                0                0 pps
```

show pfe statistics traffic egress-queues

Syntax show pfe statistics traffic egress-queues <fpc *fpc-slot*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description (On J-EX8200 switches only) Display count of multidestination packets dropped on egress ports when the egress queues are oversubscribed due to multidestination traffic.

NOTE: Multidestination packets include unknown unicast, broadcast, and multicast packets.

Options none—Displays count of packets dropped on egress ports of all physical interfaces (line cards) when egress queues are oversubscribed due to multidestination traffic.

fpc *fpc-slot*—(Optional) Displays count of packets dropped on egress ports of the physical interface (line card) referred to by the slot number.

NOTE: On a J-EX8200 switch, the FPC slot number is the slot number for the line card. Possible values are 0 through 7 on the J-EX8208 switch and 0 through 15 on the J-EX8216 switch.

Required Privilege Level view

Related Documentation

- **show pfe statistics traffic cpu on page 2997**
- **show pfe statistics traffic multicast on page 3003**
- **show interfaces queue on page 1016**
- Monitoring Interface Status and Traffic on page 931
- Understanding Junos OS CoS Components for J-EX Series Switches on page 2862

List of Sample Output **show pfe statistics traffic egress-queues fpc 4 (J-EX8208 Switch) on page 3001**

Output Fields Table 393 on page 3001 lists the output fields for the **show pfe statistics traffic egress-queues** command. Output fields are listed in the approximate order in which they appear.

Table 393: show pfe statistics traffic egress-queues Output Fields

Field Name	Field Description
Tail-dropped packets	Number of arriving packets dropped because the output queue buffers are full.

show pfe statistics traffic egress-queues

```
user@switch> show pfe statistics traffic egress-queues fpc 4
```

fpc 4 (J-EX8208 Switch)

```
Tail-dropped packets :                    0
```

show pfe statistics traffic multicast

Syntax show pfe statistics traffic multicast <fpc *fpc-slot*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description (On J-EX8200 switches only) Display class-of-service (CoS) queue information for multidestination traffic on a physical interface (line card).

NOTE: Multidestination packets include unknown unicast, broadcast, and multicast packets.

NOTE: To view statistical information for unicast traffic, use the show interfaces queue command.

Options fpc *fpc-slot*—(Optional) Displays class-of-service (CoS) queue information for multidestination traffic on the physical interface (line card) referred to by the slot number.

NOTE: On a J-EX8200 switch, the FPC slot number is the slot number for the line card. Possible values are 0 through 7 on the J-EX8208 switch and 0 through 15 on the J-EX8216 switch.

Required Privilege Level view

Related Documentation

- **show pfe statistics traffic cpu on page 2997**
- **show pfe statistics traffic egress-queues on page 3001**
- **show interfaces queue on page 1016**
- Monitoring Interface Status and Traffic on page 931
- Understanding Junos OS CoS Components for J-EX Series Switches on page 2862

List of Sample Output **show pfe statistics traffic multicast fpc 0 (J-EX8208 Switch) on page 3004**

Output Fields Table 394 on page 3003 lists the output fields for the **show pfe statistics traffic multicast** command. Output fields are listed in the approximate order in which they appear.

Table 394: show pfe statistics traffic multicast Output Fields

Field Name	Field Description
Queue	CoS queue number.

Table 394: show pfe statistics traffic multicast Output Fields *(continued)*

Field Name	Field Description
Forwarding classes	Forwarding class name.
Queued Packets	Number of packets queued to this queue.
Queued Bytes	Number of bytes queued to this queue.
Packets	Number of packets transmitted by this queue.
Bytes	Number of bytes transmitted by this queue.
Tail-dropped packets	Count of packets dropped at the tail end of the queue because of lack of buffer space.
RED-dropped packets	Number of packets dropped because of Random Early Discard (RED): • **Low**—Number of low-loss priority packets dropped because of RED. • **High**—Number of high-loss priority packets dropped because of RED.
RED-dropped bytes	Number of bytes dropped because of Random Early Discard (RED): • **Low**—Number of low-loss priority bytes dropped because of RED. • **High**—Number of high-loss priority bytes dropped because of RED.
Multicast Replication Engine-dropped packets	Egress packets dropped by the PFE because none of the ports on the physical interface are needed to forward the packet.

show pfe statistics traffic multicast fpc 0 (J-EX8208 Switch)

```
user@switch> show pfe statistics traffic multicast fpc 0

Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets              : Not Available
    Bytes                : Not Available
    Packets              :                      0                      0 pps
    Bytes                :                      0                      0 bps
    Tail-dropped packets :                      0
    RED-dropped bytes    :                      0                      0 bps
     Low                 :                      0                      0 bps
     High                :                      0                      0 bps
    RED-dropped packets  :                      0                      0 pps
     Low                 :                      0                      0 pps
     High                :                      0                      0 pps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets              : Not Available
    Bytes                : Not Available
    Packets              :                      0                      0 pps
    Bytes                :                      0                      0 bps
    Tail-dropped packets :                      0
    RED-dropped bytes    :                      0                      0 bps
     Low                 :                      0                      0 bps
     High                :                      0                      0 bps
    RED-dropped packets  :                      0                      0 pps
     Low                 :                      0                      0 pps
```

```
     High                 :                       0                       0 pps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets              : Not Available
    Bytes                : Not Available
    Packets              :                       0                       0 pps
    Bytes                :                       0                       0 bps
    Tail-dropped packets :                       0
    RED-dropped bytes    :                       0                       0 bps
     Low                 :                       0                       0 bps
     High                :                       0                       0 bps
    RED-dropped packets  :                       0                       0 pps
     Low                 :                       0                       0 pps
     High                :                       0                       0 pps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets              : Not Available
    Bytes                : Not Available
    Packets              :                       0                       0 pps
    Bytes                :                       0                       0 bps
    Tail-dropped packets :                       0
    RED-dropped bytes    :                       0                       0 bps
     Low                 :                       0                       0 bps
     High                :                       0                       0 bps
    RED-dropped packets  :                       0                       0 pps
     Low                 :                       0                       0 pps
     High                :                       0                       0 pps
Queue: 4
    Packets              : Not Available
    Bytes                : Not Available
    Packets              :                       0                       0 pps
    Bytes                :                       0                       0 bps
    Tail-dropped packets :                       0
    RED-dropped bytes    :                       0                       0 bps
     Low                 :                       0                       0 bps
     High                :                       0                       0 bps
    RED-dropped packets  :                       0                       0 pps
     Low                 :                       0                       0 pps
     High                :                       0                       0 pps
Queue: 5
    Packets              : Not Available
    Bytes                : Not Available
    Packets              :                       0                       0 pps
    Bytes                :                       0                       0 bps
    Tail-dropped packets :                       0
    RED-dropped bytes    :                       0                       0 bps
     Low                 :                       0                       0 bps
     High                :                       0                       0 bps
    RED-dropped packets  :                       0                       0 pps
     Low                 :                       0                       0 pps
     High                :                       0                       0 pps
Queue: 6
    Packets              : Not Available
    Bytes                : Not Available
    Packets              :                       0                       0 pps
    Bytes                :                       0                       0 bps
    Tail-dropped packets :                       0
    RED-dropped bytes    :                       0                       0 bps
     Low                 :                       0                       0 bps
     High                :                       0                       0 bps
    RED-dropped packets  :                       0                       0 pps
```

```
     Low                 :                       0                       0 pps
     High                :                       0                       0 pps
Queue: 7
    Packets              : Not Available
    Bytes                : Not Available
    Packets              :                       0                       0 pps
    Bytes                :                       0                       0 bps
    Tail-dropped packets :                       0
    RED-dropped bytes    :                       0                       0 bps
     Low                 :                       0                       0 bps
     High                :                       0                       0 bps
    RED-dropped packets  :                       0                       0 pps
     Low                 :                       0                       0 pps
     High                :                       0                       0 pps
Multicast Replication Engine-dropped packets : 0 pps
```

PART 22

Power over Ethernet

CHAPTER 113

Power over Ethernet (PoE)—Overview

PoE and J-EX Series Switches Overview

Power over Ethernet (PoE) permits electric power, along with data, to be passed over a copper Ethernet LAN cable. Powered devices, such as voice over IP (VoIP) telephones, wireless access points, video cameras, and point-of-sale devices, that support PoE can receive power safely from the same access ports that are used to connect personal computers to the network.

This topic describes PoE on J-EX Series Switches. It covers:

PoE

PoE was first defined in the IEEE 802.3af standard. In this standard, the amount of power that can be supplied to a powered device is limited to 15.4 W.

Whether a J-EX Series switch supports PoE depends on the switch model. Consult your switch hardware guide for information on PoE support.

PoE Power Management

Switches that have PoE ports have a PoE controller that keeps track of the PoE power consumption on the switch and allocates power to the PoE ports. The following factors determine how the PoE controller allocates power to the PoE ports:

PoE Power Budget

The PoE controller allocates power to the PoE ports from a set PoE power budget. The PoE power budget varies according to switch model and, for switches that support power supplies of different capacities, the capacity of the installed power supply.

In switches that support power supplies of different capacities, if you change your existing power supply to a lower-capacity power supply, the PoE power budget might no longer be sufficient to power all the PoE ports on the switch. If your switch supports redundant power supplies and you have installed power supplies of different capacities, the PoE power budget is based on the wattage of the lower-capacity power supply. The number of PoE ports on the switch cannot be increased by installing a larger power supply.

You can display the PoE power budget for your switch by using the **show poe controller** command.

Power Management Mode

J-EX Series switches support two power management modes: class and static. The mode you configure for your switch determines how the maximum power for a PoE interface is derived and how power is allocated to the PoE interfaces:

- Class mode—In this mode, the maximum power for an interface is determined by the class of connected powered device. Table 395 on page 3010 lists the classes of powered devices and associated power levels.

Table 395: Class of Powered Device and Power Levels

Standard	Class	Maximum Power Delivered by PoE Port	Power Range of Powered Device
IEEE 802.3af (PoE)	0	15.4 W	0.44 through 12.95 W
	1	4.0 W	0.44 through 3.84 W
	2	7.0 W	3.84 through 6.49 W
	3	15.4 W	6.49 through 12.95 W

  The powered device communicates to the PoE controller which class it belongs to when it is connected. The PoE controller then allocates to the interface the maximum power required by the class (see Table 395 on page 3010). It does not allocate power to an interface until a powered device is connected. Class 0 is the default class for powered devices that do not provide class information.

- Static mode—In this mode, you specify the maximum power for each PoE interface. The PoE controller then allocates this amount of power to the interface from its total budget. For example, if you specify a maximum value of 8.0 W for **ge-/0/0/3**, the PoE controller allocates 8.0 W out of its total power budget for the interface. This amount is allocated to the interface whether or not a powered device is connected to the interface or whether the connected powered device uses less power than 8.0 W.

  For switches that support IEEE 802.3af (PoE), the maximum power permitted on any interface is 15.4 W. This wattage guarantees that, after line loss, the powered device receives 12.95 W, which is the maximum required by 802.3af-compliant powered devices.

In both class and static mode, if the power consumption of a powered device exceeds the maximum power allocated to the interface, power to the interface is turned off.

PoE Interface Power Priority

You can configure a PoE interface to have either a high or low power priority. The power priority determines which interfaces receive power if PoE power demands are greater than the PoE power budget. If the total power allocated for all interfaces exceeds the switch budget, the lower priority interfaces are turned off and the power allocated to those interfaces drops to 0. Thus you should set interfaces that connect powered devices such as security cameras and emergency phones to high priority.

Among PoE interfaces that have the same assigned priority, power priority is determined by the port number, with lower-numbered ports having higher priority.

PoE Configuration and Monitoring

The factory default configuration enables PoE on switches that support PoE. By default, the power management mode is class, and the power priority of all interfaces is low.

If the default configuration meets your needs, you do not need to configure PoE before you connect powered devices to the switch.

To monitor the powered devices and to manage PoE power consumption, you can use the CLI or J-Web interface to display the current power consumption of the PoE ports. You can also enable the monitoring of power consumption on a port over time and then view the collected records using the CLI or the J-Web interface.

Related Documentation

- Example: Configuring PoE Interfaces on a J-EX Series Switch on page 3013
- Example: Configuring PoE Interfaces with Different Priorities on a J-EX Series Switch on page 3015

CHAPTER 114

Examples: PoE Configuration

Example: Configuring PoE Interfaces on a J-EX Series Switch

Power over Ethernet (PoE) ports supply electric power over the same ports that are used to connect network devices and allow you to plug in devices that require both network connectivity and electric power, such as voice over IP (VoIP) phones, wireless access points, and some IP cameras.

You do not need to configure PoE unless you wish to modify the default values or disable PoE on a specific interface.

This example describes a default configuration of PoE interfaces on a J-EX Series switch:

Requirements

This example uses the following software and hardware components:

- One J-EX Series switch that supports PoE

Before you configure PoE, be sure you have:

- Performed the initial switch configuration. See “Connecting and Configuring a J-EX Series Switch (CLI Procedure)” on page 161 or “Connecting and Configuring a J-EX Series Switch (J-Web Procedure)” on page 163 for details.

Overview and Topology

The topology used in this example consists of a switch that has 24 ports. Eight of the ports support PoE (IEEE 802.3af), which means they provide both network connectivity and electric power for powered devices such as VoIP telephones, wireless access points,

and IP security cameras that require 12.95 W or less. The remaining 16 ports provide only network connectivity. You use the standard ports to connect devices that have their own power sources, such as desktop and laptop computers, printers, and servers. Table 396 on page 3014 details the topology used in this configuration example.

Table 396: Components of the PoE Configuration Topology

Property	Settings
Switch hardware	J-EX Series switch with 24 Gigabit Ethernet ports: 8 PoE interfaces (**ge-0/0/0** through **ge-0/0/7**) and 16 non-PoE interfaces (**ge-0/0/8** through **ge-0/0/23**)
VLAN name	**default**
Connection to a wireless access point (requires PoE)	**ge-0/0/0**
Connections to Avaya IP telephones with integrated hubs that allow phone and desktop PC to connect to a single port (requires PoE)	**ge-0/0/1** through **ge-0/0/7**
Direct connections to desktop PCs, file servers, integrated printer/fax/copier machines (no PoE required)	**ge-0/0/8** through **ge-0/0/20**
Unused ports (for future expansion)	**ge-0/0/21** through **ge-0/0/23**

Configuration

To enable the default PoE configuration on the switch, perform these tasks:

CLI Quick Configuration

To quickly enable the default configuration on the switch:

Simply connect the powered devices to the PoE ports.

Step-by-Step Procedure

To use the PoE interfaces with default values:

1. Make sure the switch is powered on.
2. Connect the wireless access point to interface **ge-0/0/0**.
3. Connect the Avaya phones to interfaces **ge-0/0/1** through **ge-0/0/7**.

Verification

To verify that PoE interfaces have been created and are operational, perform this task:

- Verifying That the PoE Interfaces Have Been Created on page 3014

Verifying That the PoE Interfaces Have Been Created

Purpose

Verify that the PoE interfaces have been created on the switch.

Action

List all the PoE interfaces configured on the switch:

```
user@switch> show poe interface
```

```
Interface Admin status Oper status Max power Priority Power consumption Class
 ge-0/0/0 Enabled      ON          15.4W     Low      7.9W                0
 ge-0/0/1 Enabled      ON          15.4W     Low      3.2W                2
 ge-0/0/2 Enabled      ON          15.4W     Low      3.2W                2
 ge-0/0/3 Enabled      ON          15.4W     Low      3.2W                2
 ge-0/0/4 Enabled      ON          15.4W     Low      3.2W                2
 ge-0/0/5 Enabled      ON          15.4W     Low      3.2W                2
 ge-0/0/6 Enabled      ON          15.4W     Low      3.2W                2
 ge-0/0/7 Enabled      ON          15.4W     Low      3.2W                2
```

Meaning The **show poe interface** command lists PoE interfaces configured on the switch, with their status, priority, power consumption, and class. This output shows that eight interfaces have been created with default values and are consuming power at the expected rates.

Related Documentation

- Example: Configuring PoE Interfaces with Different Priorities on a J-EX Series Switch on page 3015
- Configuring PoE (CLI Procedure) on page 3021
- Troubleshooting PoE Interfaces on page 3031

Example: Configuring PoE Interfaces with Different Priorities on a J-EX Series Switch

Power over Ethernet (PoE) ports supply electric power over the same ports that are used to connect network devices. These ports allow you to plug in devices that need both network connectivity and electric power, such as voice over IP (VoIP) phones, wireless access points, and some IP cameras.

By default, PoE ports on J-EX Series switches are set to low power priority. You can configure a PoE port to have a high power priority setting. If a situation arises where there is not sufficient power for all the PoE ports, the available power is directed to the higher priority ports, while power to the lower priority ports is shut down as needed. Thus you should set ports that connect to security cameras, emergency phones, and other high priority powered devices to high priority.

This example describes how to configure a few high priority PoE interfaces.

Requirements

This example uses the following software and hardware components:

- One J-EX Series switch that supports PoE

Before you configure PoE, be sure you have:

- Performed the initial switch configuration. See “Connecting and Configuring a J-EX Series Switch (CLI Procedure)” on page 161 or “Connecting and Configuring a J-EX Series Switch (J-Web Procedure)” on page 163 for details.

Overview and Topology

The topology used in this example consists of a switch that has 24 ports. Eight of the ports support PoE (IEEE 802.3af), which means they provide both network connectivity and electric power for powered devices such as VoIP telephones, wireless access points, and IP security cameras that require 12.95 W or less. The remaining 16 ports provide only network connectivity. You use the standard ports to connect devices that have their own power sources, such as desktop and laptop computers, printers, and servers. Table 397 on page 3016 details the topology used in this configuration example.

Table 397: Components of the PoE Configuration Topology

Property	Settings
Switch hardware	Switch with 24 Gigabit Ethernet ports: 8 PoE interfaces (**ge-0/0/0** through **ge-0/0/7**) and 16 non-PoE interfaces (**ge-0/0/8** through **ge-0/0/23**)
VLAN name	**default**
Connection to a wireless access point (requires PoE)	**ge-0/0/0**
Security IP Cameras (require PoE)	**ge-0/0/1** and **ge-0/0/2 high**
Emergency VoIP phone (requires PoE)	**ge-0/0/3 high**
VoIP phone in Executive Office (requires PoE)	**ge-0/0/4 high**
Other VoIP phones (require PoE)	**ge-0/0/5** through **ge-0/0/7**
Direct connections to desktop PCs, file servers, integrated printer/fax/copier machines (no PoE required)	**ge-0/0/8** through **ge-0/0/20**
Unused ports (for future expansion)	**ge-0/0/21** through **ge-0/0/23**

Configuration

To configure PoE interfaces:

CLI Quick Configuration

By default, PoE interfaces are created for all PoE ports and PoE is enabled. The default priority for PoE interfaces is **low**.

To quickly set some interfaces to high priority and to include descriptions of the interfaces, copy the following commands and paste them into the switch terminal window:

```
[edit]
set poe interface ge-0/0/1 priority high telemetries
set poe interface ge-0/0/2 priority high telemetries
set poe interface ge-0/0/3 priority high telemetries
set poe interface ge-0/0/4 priority high telemetries
```

```
set interfaces ge-0/0/0 description "wireless access point"
set interfaces ge-0/0/1 description "security camera front door"
set interfaces ge-0/0/2 description "security camera back door"
set interfaces ge-0/0/3 description "emergency phone"
set interfaces ge-0/0/4 description "Executive Office VoIP phone"
set interfaces ge-0/0/5 description "staff VoIP phone"
set interfaces ge-0/0/6 description "staff VoIP phone"
set interfaces ge-0/0/7 description "staff VoIP phone"
```

Step-by-Step Procedure

To configure PoE interfaces with different priorities:

1. Set the interfaces connected to high priority powered devices to high priority. Include the **telemetries** statement for the high priority interfaces, thus enabling the logging of power consumption on those interfaces:

   ```
   [edit poe]
   user@switch# set interface ge-0/0/1 priority high telemetries
   user@switch# set interface ge-0/0/2 priority high telemetries
   user@switch# set interface ge-0/0/3 priority high telemetries
   user@switch# set interface ge-0/0/4 priority high telemetries
   ```

2. Provide descriptions for the PoE interfaces:

   ```
   [edit interfaces]
   user@switch# set ge-0/0/0 description "wireless access point"
   user@switch# set ge-0/0/1 description "security camera front door"
   user@switch# set ge-0/0/2 description "security camera back door"
   user@switch# set ge-0/0/3 description "emergency phone"
   user@switch# set ge-0/0/4 description "Executive Office VoIP phone"
   user@switch# set ge-0/0/5 description "staff VoIP phone"
   user@switch# set ge-0/0/6 description "staff VoIP phone"
   user@switch# set ge-0/0/7 description "staff VoIP phone"
   ```

3. Connect the wireless access point to interface **ge-0/0/0.** This interface uses the default PoE settings.

4. Connect the two security cameras to interfaces **ge-0/0/1** and **ge-0/0/2**. These interfaces are set to high priority with telemetries enabled.

5. Connect the emergency VoIP phone to interface **ge-0/0/3**. This interface is set to high priority with telemetries enabled.

6. Connect the Executive Office VoIP phone to interface **ge-0/0/4**. This interface is set to high priority with telemetries enabled.

7. Connect the staff VoIP phones to **ge-0/0/5**, **ge-0/0/6**, and **ge-0/0/7**. These interfaces use the default PoE settings.

Results

Check the results of the configuration:

```
[edit]
user@switch# show
interfaces {
  ge-0/0/0 {
    description "wireless access point";
    unit 0 {
      family ethernet-switching;
```

```
    }
  }
  ge-0/0/1 {
    description "security camera front door";
    unit 0 {
      family ethernet-switching;
    }
  }
  ge-0/0/2 {
    description "security camera back door";
    unit 0 {
      family ethernet-switching;
    }
  }
  ge-0/0/3 {
    description "emergency phone";
    unit 0 {
      family ethernet-switching;
    }
  }
  ge-0/0/4 {
    description "Executive Office VoIP phone";
    unit 0 {
      family ethernet-switching;
    }
  }
  ge-0/0/5 {
    description "staff VoIP phone";
    unit 0 {
      family ethernet-switching;
    }
  }
  ge-0/0/6 {
    description "staff VoIP phone";
    unit 0 {
      family ethernet-switching;
    }
  }
  ge-0/0/7 {
    description "staff VoIP phone";
    unit 0 {
      family ethernet-switching;
    }
  }
}
poe {
  interface all;
  interface ge-0/0/1 {
    priority high;
    telemetries;
  }
  interface ge-0/0/2 {
    priority high;
    telemetries;
  }
  interface ge-0/0/3 {
```

```
    priority high;
    telemetries;
  }
  interface ge-0/0/4 {
    priority high;
    telemetries;
  }
}
```

Verification

To verify that PoE interfaces have been created and are operational, perform the following tasks:

- Verifying That the PoE Interfaces Have Been Created with the Correct Priorities on page 3019

Verifying That the PoE Interfaces Have Been Created with the Correct Priorities

Purpose Verify that the PoE interfaces on the switch are now set to the correct priority settings.

Action List all the PoE interfaces configured on the switch:

```
user@switch> show poe interface
Interface Admin status Oper status Max power Priority Power consumption Class
 ge-0/0/0 Enabled      ON          15.4W     Low      7.9W                0
 ge-0/0/1 Enabled      ON          15.4W     High     4.8W                0
 ge-0/0/2 Enabled      ON          15.4W     High     4.8W                0
 ge-0/0/3 Enabled      ON          15.4W     High     3.3W                2
 ge-0/0/4 Enabled      ON          15.4W     High     4.7W                2
 ge-0/0/5 Enabled      ON          15.4W     Low      3.2W                2
 ge-0/0/6 Enabled      ON          15.4W     Low      3.3W                2
 ge-0/0/7 Enabled      ON          15.4W     Low      3.3W                2
```

Meaning The **show poe interface** command lists PoE interfaces configured on the switch, with their status, priority, power consumption, and class. This output shows that eight PoE interfaces are enabled. Interfaces **ge-0/0/1** through **ge-0/0/4** are configured as priority **high**. The remaining PoE interfaces are configured with the default priority value of **low**.

Related Documentation

- Example: Configuring PoE Interfaces on a J-EX Series Switch on page 3013
- Configuring PoE (CLI Procedure) on page 3021
- Troubleshooting PoE Interfaces on page 3031

CHAPTER 115

Configuring PoE

Configuring PoE (CLI Procedure)

Power over Ethernet (PoE) ports supply electric power over the same ports that are used to connect network devices. These ports allow you to plug in devices that require both network connectivity and electric power, such as voice over IP (VoIP) phones, wireless access points, and some IP cameras.

For J-EX Series switches that support PoE ports, the factory default configuration enables PoE on the PoE-capable ports, with default settings in effect. You might not have to do any additional configuration if the default settings work for you. Table 398 on page 3021 shows the PoE options and their default settings for the switch as a whole and for the PoE interfaces.

Table 398: PoE Configurable Options and Default Settings

Option	Default	Description
Switch Options		
guard-band	**0** W	Reserves up to 19 W out of the PoE power budget to be used in the case of a spike in PoE power consumption.
management	**class**	Sets the PoE power management mode for the switch: • **static**—The maximum power delivered by an interface is determined by the class of the connected powered device. No power is allocated to the interface until a powered device is connected. • **class**—The maximum power delivered by an interface is statically configured and independent of the class of the connected powered device. The maximum power is allocated to the interface even if a powered device is not connected
notification-control	Not included in default configuration	When included in the configuration, enables PoE traps.
Interface Options		

Table 398: PoE Configurable Options and Default Settings *(continued)*

Option	Default	Description
priority	**low**	Sets an interface's power priority to either **low** or **high**. If power is insufficient for all PoE interfaces, the low priority interfaces are shut down before the high priority interfaces. Among interfaces that have the same assigned priority, the power priority is determined by port number—with lower- numbered ports having higher priority.
telemetries	Not included in default configuration	When included in the configuration, enables the logging of power consumption records on an interface. Logging occurs every five minutes for one hour unless you specify a different **interval** or **duration.**
maximum-power	**15.4** W (for switches that support IEEE 802.3af) **30.0** W (for switches that support IEEE 802.3at)	Sets the maximum power that can be delivered by a PoE interface. The maximum power allowed is the same as the default—either 15.4 W for switches that do not support IEEE 802.3at or 30.0 W for switches that do support it. This setting is ignored if the power management mode is **class**.
disable	Not included in default configuration	When included in the configuration, disables PoE on the interface. The interface maintains network connectivity but no longer supplies power to a connected powered device. Power is not allocated to the interface.

To configure PoE:

1. To change power management mode from the default class mode to static mode:

   ```
   [edit poe]
   user@switch# set management static
   ```

2. To reserve a specified wattage of power in case of a spike in PoE consumption:

   ```
   [edit poe]
   user@switch# set guard-band 15
   ```

3. To configure a number of interfaces with the same settings (for example, to enable telemetry collection on all interfaces):

   ```
   [edit poe]
   user@switch# set interface all telemetries
   ```

4. To configure individual interfaces with different settings:

   ```
   [edit poe]
   user@switch# set interface ge-0/0/0 priority high telemetries duration 24

   [edit poe]
   user@switch# set interface ge-0/0/1

   [edit poe]
   user@switch# set interface ge-0/0/7 disable
   ```

   When you configure an individual interface, its configuration overrides any settings you configure with the **set poe interface all** command. For example, **ge-0/0/1** in the example above retains the default settings regardless of any settings configured with the **set poe interface all** command.

Related Documentation

- Configuring PoE (J-Web Procedure) on page 3023
- Example: Configuring PoE Interfaces on a J-EX Series Switch on page 3013
- Example: Configuring PoE Interfaces with Different Priorities on a J-EX Series Switch on page 3015
- Verifying PoE Configuration and Status (CLI Procedure) on page 3028
- PoE and J-EX Series Switches Overview on page 3009

Configuring PoE (J-Web Procedure)

Power over Ethernet (PoE) ports supply electric power over the same ports that are used to connect network devices to J-EX Series switches. These ports allow you to plug in devices that require both network connectivity and electric power, such as VoIP phones, wireless access points, and some IP cameras. Using the Power over Ethernet (PoE) Configuration page in the J-Web interface, you can modify the settings of all interfaces that are PoE-enabled.

To configure PoE:

1. Select **Configure** > **Power over Ethernet**.

   The page displays a list of all interfaces except uplink ports. Specific operational details about an interface are displayed in the Details section of the page. The details include the PoE Operational Status and Port class.

   NOTE: After you make changes to the configuration in this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select Commit Options > Commit. See "Using the Commit Options to Commit Configuration Changes (J-Web Procedure)" on page 334 for details about all commit options.

2. Click one:

   - **Edit**—Changes PoE settings for the selected port as described in Table 399 on page 3023.
   - **System Settings**—Modifies general PoE settings as described in Table 400 on page 3024.

Table 399: PoE Edit Settings

Field	Description	Your Action
Enable PoE	Specifies that PoE is enabled on the interface.	Select this option to enable PoE on the interface.
Priority	Lists the power priority (Low or High) configured on ports enabled for PoE.	Set the priority as **High** or **Low**.

Table 399: PoE Edit Settings *(continued)*

Field	Description	Your Action
Maximum Power	Specifies the maximum PoE wattage available to provision active PoE ports on the switch.	Select a value in watts. If no value is specified, the default is 15.4.

Table 400: System Settings

Field	Description	Your Action
PoE Management	Specifies the power management mode. The options are: **static** and **class**. NOTE: When the power management mode is set to **class**, the maximum power value is overridden by the maximum power value of the class of power device that is connected to the switch on the PoE port.	By default the power management mode is **static**. Select **class** to change the power management mode.
Guard Band (watts)	Specifies the band to control power availability on the switch.	Enter a value to set the guard band value in watts. The default value is 0.

Related Documentation

- Configuring PoE (CLI Procedure) on page 3021
- Example: Configuring PoE Interfaces on a J-EX Series Switch on page 3013
- Example: Configuring PoE Interfaces with Different Priorities on a J-EX Series Switch on page 3015
- Monitoring PoE on page 3025
- PoE and J-EX Series Switches Overview on page 3009

CHAPTER 116

Verifying PoE Configuration

Monitoring PoE

Purpose Use the monitoring functionality to view real-time data of the power consumed by each PoE interface, and to enable and configure telemetry values. When telemetry is enabled, the software measures the power consumed by each interface and stores the data for future reference.

Action To monitor PoE using the J-Web interface, select **Monitor > Power over Ethernet**.

To monitor PoE power consumption with CLI commands in the CLI Terminal in the J-Web interface:

1. Select **Troubleshoot > CLI Terminal**.
2. Type a CLI command:
   - **show poe controller**
   - **show poe interface**
   - **show poe telemetries interface**

For detailed information about using these CLI commands to monitor PoE power consumption, see “Monitoring PoE Power Consumption (CLI Procedure)” on page 3026.

Meaning In the J-Web interface the PoE Monitoring screen is divided into two parts. The top half of the screen displays real-time data of the power consumed by each interface and a list of ports that utilize maximum power.

Select a particular interface to view a graph of the power consumed by the selected interface.

The bottom half of the screen displays telemetry information for interfaces. The Telemetry Status field displays whether telemetry has been enabled on the interface. Click the **Show Graph** button to view a graph of the telemetries. The graph can be based on power

or voltage. To modify telemetry values, click **Edit**. Specify Interval in minutes, Duration in hours, and select **Log Telemetries** to enable telemetry on the selected interface.

Related Documentation

- Configuring PoE (CLI Procedure) on page 3021
- Configuring PoE (J-Web Procedure) on page 3023
- Example: Configuring PoE Interfaces on a J-EX Series Switch on page 3013
- Example: Configuring PoE Interfaces with Different Priorities on a J-EX Series Switch on page 3015
- Monitoring PoE Power Consumption (CLI Procedure) on page 3026
- Verifying PoE Configuration and Status (CLI Procedure) on page 3028

Monitoring PoE Power Consumption (CLI Procedure)

You can monitor Power over Ethernet (PoE) power consumption, both for the switch as a whole and for individual PoE interfaces.

This topic describes how to monitor:

PoE Power Consumption for the Switch

Purpose Determine the current PoE power consumption for the switch as a whole.

Action Enter the following command:

```
user@switch> show poe controller
Controller  Maximum   Power              Guard band  Management
index       power     consumption
   0        130 W       65W               15W        Static
```

Meaning At the time the command was executed, the PoE interfaces on the switch were consuming 65 W out of the switch PoE power budget of 130 W.

Current Power Consumption for PoE Interfaces

Purpose Determine the current power consumption for individual PoE interfaces.

Action To monitor the power consumption of all PoE interfaces on the switch, use the following command:

```
user@switch> show poe interface
Interface Admin status Oper status Max power Priority Power consumption Class
 ge-0/0/0 Enabled       ON          15.4W     Low      7.4W                0
 ge-0/0/1 Enabled       ON          15.4W     High     12.0W               0
 ge-0/0/2 Enabled       ON          15.4W     Low      12.4W               0
 ge-0/0/3 Enabled       ON          7.0W      Low      5.3W                2
```

```
ge-0/0/4 Enabled       ON         4.0W      Low      4.0W              1
ge-0/0/5 Disabled     Disabled    0.0W      Low      0.0W              0
ge-0/0/6 Enabled      OFF         15.4W     Low      0.0W              0
ge-0/0/7 Disabled     Disabled    0.0W      Low      0.0W              0
```

To monitor the power consumption of an individual PoE interface, use the following command:

```
user@switch> show poe interface ge-0/0/3
PoE interface status:
PoE interface                :  ge-0/0/3
Administrative status        : Enabled
Operational status           :  ON
Power limit on the interface : 7.0W
Priority                     : Low
Power consumed               : 5.3W
Class of power device        :   2
```

Meaning Using interface **ge-0/0/3** as an example, the powered device connected to the interface was consuming 5.3 W at the time the command was executed.

Power Consumption for PoE Interfaces over Time

Purpose Monitor the power consumption of a PoE interface over a period of time. The records collected remain available for future viewing.

You can specify the intervals at which power consumption data is collected, from once every minute to once every 30 minutes. The default is once every 5 minutes. You can also specify the duration over which the records are collected, from 1 hour (default) to 24 hours.

Action To collect historical records of PoE interface power consumption and display those records:

1. Add the **telemetries** statement to the PoE interface configuration:

   ```
   [edit]
   user@switch# set poe interface ge-0/0/5 telemetries interval 10
   ```

   When you commit the configuration, record collection begins.

2. Display the collected records:

   ```
   user@switch> show poe telemetries interface ge-0/0/5 all
   Sl No    Timestamp                 Power    Voltage
      1     03-19-2010 13:00:07 UTC  3.9W     50.9V
      2     03-19-2010 12:50:07 UTC  3.9W     50.9V
      3     03-19-2010 12:40:07 UTC  3.9W     50.9V
      4     03-19-2010 12:30:07 UTC  3.9W     50.9V
      5     03-19-2010 12:20:07 UTC  3.9W     50.9V
      6     03-19-2010 12:10:07 UTC  3.9W     50.9V
   ```

To start another session of record collection on the interface, you must commit the configuration again.

Meaning Over the hour in which the PoE power consumption data on **ge-0/0/5** was collected, the connected powered device consistently consumed 3.9 W.

Related Documentation

- Configuring PoE (CLI Procedure) on page 3021
- Example: Configuring PoE Interfaces on a J-EX Series Switch on page 3013
- Example: Configuring PoE Interfaces with Different Priorities on a J-EX Series Switch on page 3015
- Verifying PoE Configuration and Status (CLI Procedure) on page 3028

Verifying PoE Configuration and Status (CLI Procedure)

You can verify the Power over Ethernet (PoE) configuration and status on a J-EX Series switch.

This topic describes how to verify the:

- Number of PoE Ports on the Switch on page 3028
- PoE Controller Configuration and Status on page 3028
- PoE Interface Configuration and Status on page 3029
- PoE SNMP Trap Generation Status on page 3029

Number of PoE Ports on the Switch

Purpose Verify the number of PoE ports on a switch. The number of PoE ports on a switch varies according to switch model.

Action Enter the following command:

```
user@switch> show chassis hardware
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                BH0208375304      EX4200-24T
Routing Engine 0 REV 11   750-021261   BH0208375304      EX4200-24T, 8 POE
FPC 0            REV 11   750-021261   BH0208375304      EX4200-24T, 8 POE
  CPU                     BUILTIN      BUILTIN           FPC CPU
  PIC 0                   BUILTIN      BUILTIN           24x 10/100/1000 Base-T
Power Supply 0   REV 03   740-020957   AT0508285661      PS 320W AC
Fan Tray                                                 Fan Tray
```

Meaning The switch is a J-EX4200-24T model with eight PoE ports.

PoE Controller Configuration and Status

Purpose Verify the PoE controller configuration and status, such as the PoE power budget, total PoE power consumption, and power management mode.

Action Enter the following command:

```
user@switch> show poe controller
Controller  Maximum   Power             Guard band  Management
```

```
index       power     consumption
   0        130 W       43W               15W        Class
```

Meaning The switch has an overall PoE power budget of 130 W, of which 43 W were being used by the PoE ports at the time the command was executed. The Guard band field shows that 15 W is reserved out of the PoE power budget to protect against spikes in power demand. The power management mode is class.

PoE Interface Configuration and Status

Purpose Verify that PoE interfaces are enabled and set to the correct maximum power and priority settings. Also verify current operational status and power consumption.

Action To view configuration and status for all PoE interfaces, enter:

```
user@switch>  show poe interface
Interface Admin status Oper status Max power Priority Power consumption Class
 ge-0/0/0 Enabled      ON          15.4W     Low      7.9W                3
 ge-0/0/1 Enabled      ON          15.4W     High     4.8W                0
 ge-0/0/2 Enabled      ON          15.4W     High     4.8W                0
 ge-0/0/3 Enabled      ON          15.4W     High     3.3W                2
 ge-0/0/4 Disabled     Disabled    0.0W      Low      0.0W                0
 ge-0/0/5 Enabled      ON          15.4W     Low      3.2W                2
 ge-0/0/6 Enabled      ON          15.4W     Low      3.3W                2
 ge-0/0/7 Enabled      OFF         15.4W     Low      0.0W                0
```

To view configuration and status for a single PoE interface, enter:

```
user@switch> show poe interface ge-0/0/3
PoE interface status:
PoE interface                :  ge-0/0/3
Administrative status        : Enabled
Operational status           :  ON
Power limit on the interface : 15.4W
Priority                     : High
Power consumed               : 3.3W
Class of power device        :   2
```

Meaning The command output shows the status and configuration of interfaces. For example, the interface **ge-0/0/3** is administratively enabled. Its operational status is **ON**; that is, the interface is currently delivering power to a connected powered device. The maximum power the interface can deliver is 15.4 W. The interface has a high power priority. At the time the command was executed, the powered device was consuming 3.3 W. The IEEE 802.3af class of the powered device is class 2.

PoE SNMP Trap Generation Status

Purpose Verify the status of the **notification-control** option, which determines whether or not PoE SNMP traps are enabled.

Action Enter the following command:

```
user@switch> show poe notification-control
FPC slot             Notification-control-status
    0                    OFF
```

Meaning PoE SNMP traps are not enabled.

Related Documentation

- Configuring PoE (CLI Procedure) on page 3021
- Example: Configuring PoE Interfaces on a J-EX Series Switch on page 3013
- Example: Configuring PoE Interfaces with Different Priorities on a J-EX Series Switch on page 3015
- Monitoring PoE Power Consumption (CLI Procedure) on page 3026

CHAPTER 117

Troubleshooting PoE Configuration

Troubleshooting PoE Interfaces

Problem A Power over Ethernet (PoE) interface is not supplying power to the powered device.

Solution Check for the items shown in Table 401 on page 3031.

Table 401: Troubleshooting a PoE Interface

Items to Check	Explanation
Is the switch a full PoE model or a partial PoE model?	If you are using a partial PoE model, only interfaces **ge-0/0/0** through **ge-0/0/7** can function as PoE ports.
Has PoE capability been disabled for that interface?	Use the **show poe interface** command to check PoE interface status.
Is the cable properly seated in the port socket?	Check the hardware.
Has the PoE power budget been exceeded for the switch?	Use the **show poe controller** command to check the PoE power budget and consumption for the switch.
Does the powered device require more power than is available on the interface?	Use the **show poe interface** command to check the maximum power provided by the interface.
If the **telemetries** option has been enabled for the interface, check the history of power consumption.	Use the **show poe telemetries interface** command to display the history of power consumption.

Related Documentation

- Example: Configuring PoE Interfaces with Different Priorities on a J-EX Series Switch on page 3015
- Verifying PoE Configuration and Status (CLI Procedure) on page 3028
- Monitoring PoE Power Consumption (CLI Procedure) on page 3026
- Configuring PoE (CLI Procedure) on page 3021

CHAPTER 118

Configuration Statements for PoE

[edit poe] Configuration Statement Hierarchy

```
poe {
  guard-band watts;
  interface (all | interface-name) {
    disable;
    maximum-power watts;
    priority (high | low);
    telemetries {
      disable;
      duration hours;
      interval minutes;
    }
  }
  management (class | static);
  notification-control {
    fpc slot-number {
      disable;
    }
  }
}
```

Related Documentation

- Example: Configuring PoE Interfaces with Different Priorities on a J-EX Series Switch on page 3015
- Configuring PoE (CLI Procedure) on page 3021
- PoE and EX Series Switches Overview on page 3009

disable

Syntax disable;

Hierarchy Level [edit poe **interface** (all | *interface-name*)],
[edit poe **interface** (all | *interface-name*) **telemetries**],
[edit poe **notification-control fpc** *slot-number*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable a PoE interface, disable the collection of power consumption data for a PoE interface, or disable the generation of the PoE SNMP traps. The action of the **disable** statement depends on which statement it is used with:

- When used with **interface**—Disable the PoE capability of this interface. The interface operates as a standard network access interface, and power is no longer allocated to it from the PoE power budget. Although the PoE capability is disabled, the PoE configuration for the interface is retained. To re-enable the PoE capability of this interface, delete the **disable** statement from the **interface** entry in the configuration.
- When used with **telemetries**—Disable the collection of PoE power consumption records for this port. Any previously collected records are deleted. However, the **telemetries** configuration is retained, including the values for **interval** and **duration**. To re-enable record collection, delete the **disable** statement from the **telemetries** entry in the configuration.
- When used with **notification-control**—Disable the generation of PoE traps. To re-enable PoE traps, delete the **disable** statement from the **notification-control** entry in the configuration.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring PoE Interfaces with Different Priorities on a J-EX Series Switch on page 3015
- Configuring PoE (CLI Procedure) on page 3021

duration

Syntax duration *hours*;

Hierarchy Level [edit poe **interface** (all | *interface-name*) **telemetries**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Modify the duration over which data is collected when you are monitoring the power consumption of a PoE interface.

Options ***hours*** —Number of hours over which the data is to be collected.

Range: 1 through 24

Default: 1

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring PoE Interfaces with Different Priorities on a J-EX Series Switch on page 3015
- Configuring PoE (CLI Procedure) on page 3021

fpc

Syntax

```
fpc slot-number {
    disable;
}
```

Hierarchy Level [edit poe notification-control]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable the generation of PoE traps for the specified FPC.

Default PoE traps are disabled by default.

Options ***slot-number***—The FPC slot number, where ***slot- number*** is:

- **0**—On a standalone J-EX4200 switch.
- **0** through **9**—On a J-EX4200 switch in a Virtual Chassis, indicating the member ID.

The remaining statement is explained separately.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring PoE Interfaces with Different Priorities on a J-EX Series Switch on page 3015
- Configuring PoE (CLI Procedure) on page 3021

guard-band

Syntax guard-band *watts*;

Hierarchy Level [edit poe]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Reserve a specified amount of power out of the PoE power budget in case of a spike in PoE consumption.

Options ***watts*** —Amount of power to be reserved in case of a spike in PoE consumption.

Range: 0 through 19

Default: 0

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring PoE Interfaces with Different Priorities on a J-EX Series Switch on page 3015
- Configuring PoE (CLI Procedure) on page 3021

interface

Syntax

```
interface (all | interface-name) {
  disable;
  maximum-power watts;
  priority (high | low);
  telemetries {
    disable;
    duration hours;
    interval minutes;
  }
}
```

Hierarchy Level [edit poe]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify a PoE interface to be configured.

Options

all—All PoE interfaces on the switch that have not been individually configured for PoE. If a PoE interface has been individually configured, that configuration overrides any settings specified with **all**.

interface-name—Name of the specific interface being configured.

If you use the **interface** statement without any substatements, PoE is enabled on all interfaces or the specified interface with default values for the remaining statements.

The remaining statements are explained separately.

Required Privilege Level

system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring PoE Interfaces with Different Priorities on a J-EX Series Switch on page 3015
- Configuring PoE (CLI Procedure) on page 3021

interval

Syntax interval *minutes*;

Hierarchy Level [edit poe **interface** (all | *interface-name*) **telemetries**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Modify the interval at which data is collected when you are monitoring the power consumption of a PoE interface.

Options ***minutes***—Frequency of data collection.

Range: 1 through 30

Default: 5

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring PoE Interfaces with Different Priorities on a J-EX Series Switch on page 3015
- Configuring PoE (CLI Procedure) on page 3021
- Configuring PoE (J-Web Procedure) on page 3023

management

Syntax management (class | static);

Hierarchy Level [edit poe]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Designate the way that the switch's PoE controller allocates power to the PoE interfaces.

Default class

Options

- **class**—The amount of power allocated to the interface is determined by the class of the connected powered device. If no powered device is connected, no power is allocated to the interface. See "PoE and J-EX Series Switches Overview" on page 3009 for more information about classes of powered devices.
- **static**—The amount of power allocated to the interface is determined by the value of the **maximum-power** statement, not the class of the connected powered device. This amount is allocated even when a powered device is not connected to the interface, ensuring that power is available when needed.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring PoE Interfaces with Different Priorities on a J-EX Series Switch on page 3015
- Configuring PoE (CLI Procedure) on page 3021

maximum-power

Syntax maximum-power *watts*;

Hierarchy Level [edit poe **interface** (all | *interface-name*)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set the maximum amount of power that the switch can supply to the PoE port.

NOTE: Although you can set this value when PoE power management is in class mode, it does not establish the maximum power for the port. Instead, the IEEE 802.3af or IEEE 802.3at class of the connected device determines the maximum power for the port.

Options ***watts*** —The maximum number of watts that can be supplied to the port.

Range: 0.0 through 15.4 for switches that support only IEEE 802.3af and 0.0 through 30.0 for switches that also support IEEE 802.3at

Default: 15.4 for switches that support IEEE 802.3af and 30.0 for switches that support IEEE 802.3at

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring PoE Interfaces with Different Priorities on a J-EX Series Switch on page 3015
- Configuring PoE (CLI Procedure) on page 3021

notification-control

Syntax

```
notification-control {
    fpc slot-number {
        disable;
    }
}
```

Hierarchy Level [edit poe]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable or disable the generation of PoE SNMP traps. If PoE traps are enabled, an SNMP trap is sent whenever a PoE interface is enabled or disabled.

The remaining statements are explained separately.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring PoE Interfaces with Different Priorities on a J-EX Series Switch on page 3015
- Configuring PoE (CLI Procedure) on page 3021

priority

Syntax priority (low | high);

Hierarchy Level [edit poe **interface** (all | *interface-name*)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set the power priority for individual interfaces when there is insufficient power for all PoE interfaces. If the switch needs to shut down powered devices because PoE demand exceeds the PoE budget, low priority devices are shut down before high priority devices. Among interfaces that have the same assigned priority, priority is determined by port number, with lower-numbered ports having higher priority.

Default **low**

Options *value*—**high** or **low**:

- **high**—Specifies that this interface is to be treated as high priority in terms of power allocation. If the switch needs to shut down powered devices because PoE demand exceeds the PoE budget, power is not shut down on this interface until it has been shut down on all the low priority interfaces.
- **low**—Specifies that this interface is to be treated as low priority in terms of power allocation. If the switch needs to shut down powered devices because PoE demand exceeds the PoE budget, power is shut down on this interface before it is shut down on high priority interfaces.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring PoE Interfaces with Different Priorities on a J-EX Series Switch on page 3015
- Configuring PoE (CLI Procedure) on page 3021

telemetries

Syntax

```
telemetries {
  disable;
  duration hours;
  interval minutes;
}
```

Hierarchy Level [edit poe **interface** (all | *interface-name*)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable the logging of power consumption of a PoE interface over time.

If you want to log the power consumption of a PoE interface, you must explicitly specify the **telemetries** statement. When you commit the configuration, logging begins, with data being collected at the specified intervals. Logging stops at the end of the specified duration. If you did not specify the **duration** and **interval** statements, data is collected at five minute intervals for one hour.

The remaining statements are explained separately.

Default Logging of power consumption is disabled.

Required Privilege Level system—To view this statement in the configuration.
system-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring PoE Interfaces with Different Priorities on a J-EX Series Switch on page 3015
- Configuring PoE (CLI Procedure) on page 3021

CHAPTER 119

Operational Mode Commands for PoE

show poe controller

Syntax **show poe controller**

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display configuration and status of the PoE controller.

Required Privilege Level view

Related Documentation

- **show poe interface on page 3048**
- Verifying PoE Configuration and Status (CLI Procedure) on page 3028
- Monitoring PoE Power Consumption (CLI Procedure) on page 3026
- Troubleshooting PoE Interfaces on page 3031

List of Sample Output **show poe controller on page 3047**

Output Fields Table 402 on page 3046 lists the output fields for the **show poe controller** command. Output fields are listed in the approximate order in which they appear.

Table 402: show poe controller Output Fields

Field Name	Field Description
Controller index	Controller number.
Maximum power	Maximum power that the switch can provide to the PoE ports.
Power consumption	Total amount of power being used by the PoE ports at the time the command is executed.
Guard Band	Amount of power that has been placed in reserve for power demand spikes and that cannot be allocated to a PoE interface.
Management	Power management mode: either **Static** or **Class**.

show poe controller `user@switch>` **show poe controller**

```
Controller  Maximum   Power              Guard band  Management
index       power     consumption
   0        130 W      43W                15W        Class
```

show poe interface

Syntax **show poe interface**
<interface-name>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display the status of PoE interfaces.

Options none—Display status of all PoE interfaces on the switch.

interface-name—(Optional) Display the status of a specific PoE interface on the switch.

Required Privilege Level view

Related Documentation

- **show poe controller on page 3046**
- Verifying PoE Configuration and Status (CLI Procedure) on page 3028
- Monitoring PoE Power Consumption (CLI Procedure) on page 3026
- Troubleshooting PoE Interfaces on page 3031

List of Sample Output **show poe interface on page 3049**
show poe interface ge-0/0/3 on page 3049

Output Fields Table 403 on page 3048 lists the output fields for the **show poe interface** command. Output fields are listed in the approximate order in which they appear.

Table 403: show poe interface Output Fields

Field Name (All Interfaces Output)	Field Name (Single Interface Output)	Field Description
Interface	**PoE Interface**	Interface name.
Admin status	**Administrative status**	Administrative state of the PoE interface: **Enabled** or **Disabled**. If the PoE interface is disabled, it can provide network connectivity, but it cannot provide power to connected devices.
Oper status	**Operational status**	Operational state of the PoE interface: • **ON**—The interface is currently supplying power to a powered device. • **OFF**—PoE is enabled on the interface, but the interface is not currently supplying power to a powered device. • **Disabled**—PoE is disabled on the interface.
Max power	**Power limit on the interface**	Maximum power that can be provided by the interface.
Priority	**Priority**	Interface power priority: either **High** or **Low**.

Table 403: show poe interface Output Fields *(continued)*

Field Name (All Interfaces Output)	Field Name (Single Interface Output)	Field Description
Power consumption	**Power consumed**	Amount of power being used by the interface at the time the command is executed.
Class	**Class of power device**	IEEE 802.3af or IEEE 802.3at classification of the powered device. Class 0 is the default class and is used when the class of the powered device is unknown or no powered device is connected.

show poe interface

```
user@switch> show poe interface

Interface Admin status Oper status Max power Priority Power consumption Class
 ge-0/0/0 Enabled      ON          15.4W     Low      7.9W                0
 ge-0/0/1 Enabled      ON          15.4W     Low      3.2W                2
 ge-0/0/2 Enabled      ON          15.4W     Low      3.2W                2
 ge-0/0/3 Enabled      ON          15.4W     Low      3.2W                2
 ge-0/0/4 Enabled      ON          15.4W     Low      3.2W                2
 ge-0/0/5 Enabled      ON          15.4W     Low      3.2W                2
 ge-0/0/6 Enabled      ON          15.4W     Low      3.2W                2
 ge-0/0/7 Enabled      ON          15.4W     Low      3.2W                2
```

show poe interface ge-0/0/3

```
user@switch> show poe interface ge-0/0/3
PoE interface status:
PoE interface                :  ge-0/0/3
Administrative status        : Enabled
Operational status           :  ON
Power limit on the interface : 7.0W
Priority                     : Low
Power consumed               : 5.3W
Class of power device        :   2
```

show poe notification-control

Syntax **show poe notification-control**

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display the state of the PoE **notification-control** option, which enables or disables PoE SNMP traps.

Required Privilege Level view

Related Documentation

- **show poe controller on page 3046**
- **show poe interface on page 3048**
- Verifying PoE Configuration and Status (CLI Procedure) on page 3028

List of Sample Output **show poe notification-control on page 3051**

Output Fields Table 404 on page 3050 lists the output fields for the **show poe notification-control** command. Output fields are listed in the approximate order in which they appear.

Table 404: show poe notification-control Output Fields

Field Name	Field Description
FPC slot	FPC slot number.
Notification-control-status	Status of notification control: • **ON**—PoE traps are enabled. An SNMP trap is sent when a PoE interface is enabled or disabled. • **OFF**—PoE traps are disabled.

show poe notification-control

```
user@switch> show poe notification-control
FPC slot              Notification-control-status
    0                     OFF
```

show poe telemetries interface

Syntax **show poe telemetries interface** ***interface-name*** **(all | *n*)**

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display a history of power consumption on the specified interface.

Telemetries must be enabled on the interface before you can display a history of power consumption.

Options ***interface-name***—Display power consumption records for the specified PoE interface.

all—Display all power consumption records for the PoE interface.

n—Display the specified number of power consumption records for the PoE interface. The records displayed are the most recent.

Required Privilege Level view

Related Documentation

- **show poe interface on page 3048**
- **show poe controller on page 3046**
- Monitoring PoE Power Consumption (CLI Procedure) on page 3026
- Verifying PoE Configuration and Status (CLI Procedure) on page 3028
- Troubleshooting PoE Interfaces on page 3031

List of Sample Output **show poe telemetries interface (Last 10 Records) on page 3053**
show poe telemetries interface (All Records) on page 3053

Output Fields Table 405 on page 3052 lists the output fields for the **show poe telemetries interface** command. Output fields are listed in the approximate order in which they appear.

Table 405: show poe telemetries interface Output Fields

Field Name	Field Description
Sl No	Number of the record for the specified port. Record number 1 is the most recent.
Timestamp	Date and time when the power-consumption data was gathered.
Power	Amount of power provided by the specified interface at the time the data was gathered.
Voltage	Maximum voltage provided by the specified interface at the time the data was gathered.

show poe telemetries interface (Last 10 Records)

```
user@switch> show poe telemetries interface ge-0/0/0 10
Sl No    Timestamp                 Power    Voltage
   1     01-27-2008 18:19:58 UTC   15.4W     51.6V
   2     01-27-2008 18:18:58 UTC   15.4W     51.6V
   3     01-27-2008 18:17:58 UTC   15.4W     51.6V
   4     01-27-2008 18:16:58 UTC   15.4W     51.6V
   5     01-27-2008 18:15:58 UTC   15.4W     51.6V
   6     01-27-2008 18:14:58 UTC   15.4W     51.6V
   7     01-27-2008 18:13:58 UTC   15.4W     51.6V
   8     01-27-2008 18:12:57 UTC   15.4W     51.6V
   9     01-27-2008 18:11:57 UTC   15.4W     51.6V
  10     01-27-2008 18:10:57 UTC   15.4W     51.6V
```

show poe telemetries interface (All Records)

```
user@switch> show poe telemetries interface ge-0/0/0 all
Sl No    Timestamp                 Power    Voltage
   1     01-27-2008 18:19:58 UTC   15.4W     51.6V
   2     01-27-2008 18:18:58 UTC   15.4W     51.6V
   3     01-27-2008 18:17:58 UTC   15.4W     51.6V
   4     01-27-2008 18:16:58 UTC   15.4W     51.6V
   5     01-27-2008 18:15:58 UTC   15.4W     51.6V
   6     01-27-2008 18:14:58 UTC   15.4W     51.6V
   7     01-27-2008 18:13:58 UTC   15.4W     51.6V
   8     01-27-2008 18:12:57 UTC   15.4W     51.6V
   9     01-27-2008 18:11:57 UTC   15.4W     51.6V
  10     01-27-2008 18:10:57 UTC   15.4W     51.6V
  11     01-27-2008 18:09:57 UTC   15.4W     51.6V
  12     01-27-2008 18:08:57 UTC   15.4W     51.6V
  13     01-27-2008 18:07:57 UTC   15.4W     51.6V
  14     01-27-2008 18:06:57 UTC   15.4W     51.6V
  15     01-27-2008 18:05:57 UTC   15.4W     51.6V
  16     01-27-2008 18:04:56 UTC   15.4W     51.6V
  17     01-27-2008 18:03:56 UTC   15.4W     51.6V
  18     01-27-2008 18:02:56 UTC   15.4W     51.6V
  19     01-27-2008 18:01:56 UTC   15.4W     51.6V
  20     01-27-2008 18:00:56 UTC   15.4W     51.6V
  21     01-27-2008 17:59:56 UTC   15.4W     51.6V
  22     01-27-2008 17:58:56 UTC   15.4W     51.6V
  23     01-27-2008 17:57:56 UTC   15.4W     51.6V
  24     01-27-2008 17:56:55 UTC   15.4W     51.6V
  25     01-27-2008 17:55:55 UTC   15.4W     51.6V
  26     01-27-2008 17:54:55 UTC   15.4W     51.6V
  27     01-27-2008 17:53:55 UTC   15.4W     51.6V
  28     01-27-2008 17:52:55 UTC   15.4W     51.6V
  29     01-27-2008 17:51:55 UTC   15.4W     51.6V
  30     01-27-2008 17:50:55 UTC   15.4W     51.6V
  31     01-27-2008 17:49:55 UTC   15.4W     51.6V
  32     01-27-2008 17:48:55 UTC   15.4W     51.6V
  33     01-27-2008 17:47:54 UTC   15.4W     51.6V
  34     01-27-2008 17:46:54 UTC   15.4W     51.6V
  35     01-27-2008 17:45:54 UTC   15.4W     51.6V
  36     01-27-2008 17:44:54 UTC   15.4W     51.6V
  37     01-27-2008 17:43:54 UTC   15.4W     51.6V
  38     01-27-2008 17:42:54 UTC   15.4W     51.6V
  39     01-27-2008 17:41:54 UTC   15.4W     51.6V
  40     01-27-2008 17:40:54 UTC   15.4W     51.6V
  41     01-27-2008 17:39:53 UTC   15.4W     51.6V
  42     01-27-2008 17:38:53 UTC   15.4W     51.6V
  43     01-27-2008 17:37:53 UTC   15.4W     51.6V
  44     01-27-2008 17:36:53 UTC   15.4W     51.6V
```

PART 23

MPLS

CHAPTER 120

MPLS—Overview

Junos OS MPLS for J-EX Series Switches Overview

You can configure Junos OS MPLS on J-EX Series Switches to increase transport efficiency in the network. MPLS services can be used to connect various sites to a backbone network and to ensure better performance for low-latency applications such as VoIP and other business-critical functions.

Junos OS MPLS for J-EX Series switches supports:

- Layer 2 protocols
- Layer 2 VPNs
- RSVP-based label-switched paths (LSPs)
- MPLS-based circuits cross-connect (CCCs)
- IP over MPLS
- Class of service (CoS)

NOTE: MPLS configurations on J-EX Series switches are compatible with configurations on other devices running Junos OS that support MPLS and CCC.

Benefits of MPLS

MPLS has the following advantages over conventional packet forwarding:

- Packets arriving on different ports can be assigned different labels.
- A packet arriving at a particular provider edge switch may be assigned a label that is different from that of the same packet entering the network at a different provider edge switch. As a result, forwarding decisions that depend on the ingress provider edge switch can be easily made.
- Sometimes it is desirable to force a packet to follow a particular route that is explicitly chosen at or before the time the packet enters the network, rather than letting it follow the route chosen by the normal dynamic routing algorithm as the packet travels through the network. In MPLS, a label can be used to represent the route so that the packet need not carry the identity of the explicit route.

NOTE: MPLS configurations on J-EX Series switches do not support:

- LDP-based MPLS
- Routed VLAN interfaces (RVIs)
- Q-in-Q tunneling
- Aggregated Ethernet interfaces (LAGs) on CCCs
- CCCs with a beginning and ending on the same switch

Additional Benefits of MPLS and Traffic Engineering

MPLS is the packet-forwarding component of the Junos OS traffic engineering architecture. Traffic engineering provides the capabilities to do the following:

- Route primary paths around known bottlenecks or points of congestion in the network.
- Provide precise control over how traffic is rerouted when the primary path is faced with single or multiple failures.
- Provide efficient use of available aggregate bandwidth and long-haul fiber by ensuring that certain subsets of the network are not overutilized while other subsets of the network along potential alternate paths are underutilized.
- Maximize operational efficiency.
- Enhance the traffic-oriented performance characteristics of the network by minimizing packet loss, minimizing prolonged periods of congestion, and maximizing throughput.
- Enhance statistically bound performance characteristics of the network (such as loss ratio, delay variation, and transfer delay) required to support a multiservice Internet.

Related Documentation

- Understanding MPLS Label Operations on J-EX Series Switches on page 3067
- Understanding Junos OS MPLS Components for J-EX Series Switches on page 3059
- Understanding Using CoS with MPLS Networks on J-EX Series Switches on page 2880
- Example: Configuring MPLS on J-EX Series Switches on page 3071

- *Junos OS MPLS Applications Configuration Guide* at http://www.juniper.net/techpubs/software/junos/
- *Junos OS VPNs Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

Understanding Junos OS MPLS Components for J-EX Series Switches

Junos OS MPLS for J-EX Series Switches supports Layer 2 protocols and Layer 2 virtual private networks (VPNs). You can configure MPLS on J-EX Series switches to increase transport efficiency in your network. MPLS services can be used to connect various sites to a backbone network and to ensure better performance for low-latency applications such as VoIP and other business-critical functions.

This topic includes:

Provider Edge Switches

To implement MPLS on J-EX Series switches, you must configure two provider edge (PE) switches—that is, an ingress (local) PE switch and an egress (remote) PE switch.

The ingress switch (the entry point to the MPLS tunnel) receives an IP packet, analyzes it, and pushes an MPLS label onto it. This label places the packet in a forwarding equivalence class (FEC) and determines its handling and destination through the MPLS tunnel. The egress provider edge switch (the exit point from the MPLS tunnel) pops the MPLS label off the outgoing packet.

MPLS traffic is bidirectional. Therefore, each PE switch can be configured as both an ingress switch and an egress switch, depending on the direction of the traffic.

J-EX Series switches can handle only single-label MPLS packets. If a packet already has an MPLS label, the PE switch removes the label and swaps it for another MPLS label.

MPLS Protocol and Label Switched Paths

Each PE switch must be configured to support the MPLS protocol, and the MPLS stanza must include the configuration of a label-switched path (LSP) that specifies the address of the remote PE switch.

Junos OS MPLS for J-EX Series switches supports RSVP-based LSPs.

Circuit Cross-Connect for Customer-Edge Interfaces

You can configure the customer-edge interfaces of the PE switches as a circuit cross-connect (CCC), to create a transparent connection between two circuits. When you configure an interface as a CCC, the interface no longer belongs to a default VLAN. The interface becomes an MPLS tunnel — used exclusively for MPLS packets. You can

create different CCCs for different customers or for segregating different traffic streams over different MPLS tunnels.

Using CCC, you can connect the following types of circuits:

- Local interface with remote interface or VLAN
- Local VLAN with remote interface or VLAN

NOTE: To configure a VLAN circuit as a CCC, you must enable VLAN tagging and specify a VLAN ID.

MPLS on J-EX Series switches does not support the following types of CCC configurations:

- LDP-based MPLS
- Routed VLAN interfaces (RVIs)
- Q-in-Q tunneling
- Aggregated Ethernet interfaces (LAGs) on CCCs
- CCCs with a beginning and ending on the same switch

IP over MPLS For Customer-Edge Interfaces

You can configure the customer-edge interfaces of the PE switches for IP over MPLS using a Layer 3 interface and a static route from the ingress PE switch to the egress PE switch. See “Configuring MPLS on Provider Edge Switches Using IP Over MPLS (CLI Procedure)” on page 3107 for additional information.

Provider Switch

You must configure one or more provider switches as transit switches within the network to support the forwarding of MPLS packets. You can add provider switches without changing the configuration of the PE switches.

A provider switch does not analyze the packets. It refers to an MPLS label forwarding table and swaps one label for another. The new label determines the next hop along the MPLS tunnel. A provider switch cannot perform the push or pop operations.

Components Required for All Switches in the MPLS Network

You must configure the following components on both the provider edge and the provider switches:

- Routing Protocol on page 3061
- Traffic Engineering on page 3061
- MPLS Protocol on page 3061
- RSVP on page 3061

Routing Protocol

MPLS works in coordination with the interior gateway protocol (IGP). Therefore, you must configure OSPF or IS-IS as the routing protocol on the loopback interface and core interfaces of both the provider edge and provider switches.

These core interfaces can be either Gigabit Ethernet or 10-Gigabit Ethernet interfaces, and they can be configured as either individual interfaces or aggregated Ethernet interfaces.

NOTE: These core interfaces cannot be configured with VLAN tagging or a VLAN ID. When you configure them to belong to **family mpls**, they are removed from the default VLAN. They operate as an exclusive tunnel for MPLS traffic.

Traffic Engineering

Traffic engineering maps traffic flows onto an existing physical topology and provides the ability to move traffic flow away from the shortest path selected by the IGP and onto a potentially less congested physical path across a network.

Traffic engineering enables the selection of specific end-to-end paths to send given types of traffic through your network. For MPLS to work properly, you must enable traffic engineering for the specified routing protocol.

MPLS Protocol

You must enable the MPLS protocol on all switches that participate in the MPLS network and apply it to the core interfaces of both the provider edge and provider switches. You do not need to apply it to the loopback interface, because the MPLS protocol uses the framework established by the RSVP session to create LSPs. On the provider edge switches, the configuration of the MPLS protocol must also include the definition of an LSP.

RSVP

Resource Reservation Protocol (RSVP) is a signaling protocol that allocates and distributes labels throughout an MPLS network. RSVP sets up unidirectional paths between the ingress provider edge switch and the egress provider edge switch. RSVP makes the LSPs dynamic; it can detect topology changes and outages and establish new LSPs to move around a failure.

You must enable RSVP and apply it to the loopback interface and the core interface of both the provider edge and provider switches. The path message contains the configured information about the resources required for the LSP to be established.

When the egress switch receives the path message, it sends a reservation message back to the ingress switch. This reservation message is passed along from switch to switch along the same path as the original path message. Once the ingress switch receives this reservation message, an RSVP path is established.

The established LSP stays active as long as the RSVP session remains active. RSVP continues activity through the transmissions and responses to RSVP path and reservation

messages. If the messages stop for three minutes, the RSVP session terminates and the LSP is lost.

RSVP runs as a separate software process in the Junos OS and is not in the packet-forwarding path.

Family MPLS

You must configure the core interfaces used for MPLS traffic to belong to **family mpls**.

NOTE: You can enable **family mpls** on either individual interfaces or aggregated Ethernet interfaces. You cannot enable it on tagged VLAN interfaces.

Related Documentation

- Junos OS MPLS for J-EX Series Switches Overview on page 3057
- Understanding MPLS and Path Protection on J-EX Series Switches on page 3063
- Example: Configuring MPLS on J-EX Series Switches on page 3071
- Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect (CLI Procedure) on page 3111
- Configuring MPLS on Provider Edge Switches Using IP Over MPLS (CLI Procedure) on page 3107
- Configuring MPLS on Provider Switches (CLI Procedure) on page 3102
- *Junos OS MPLS Applications Configuration Guide* at http://www.juniper.net/techpubs/software/junos/
- *Junos OS VPNs Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

Understanding MPLS and Path Protection on J-EX Series Switches

Junos OS MPLS for J-EX Series Switches provides path protection to protect your MPLS network from label switched path (LSP) failures.

By default, an LSP routes itself hop-by-hop from the ingress provider edge switch through the provider switches toward the egress provider edge switch. The LSP generally follows the shortest path as dictated by the local routing table, usually taking the same path as destination-based, best-effort traffic. These paths are "soft" in nature because they automatically reroute themselves whenever a change occurs in a routing table or in the status of a node or link.

Typically, when an LSP fails, the switch immediately upstream from the failure signals the outage to the ingress provider edge switch. The ingress provider edge switch calculates a new path to the egress provider edge switch, establishes the new LSP, and then directs traffic from the failed path to the new path. This rerouting process can be time-consuming and prone to failure. For example, the outage signals to the ingress switch might get lost or the new path might take too long to come up, resulting in significant packet drops.

You can configure path protection by configuring primary and secondary paths on the ingress switch. If the primary path fails, the ingress switch immediately reroutes traffic from the failed path to the standby path, eliminating the need for the ingress switch to calculate a new route and signal a new path. For information about configuring standby LSPs, see "Configuring Path Protection in an MPLS Network (CLI Procedure)" on page 3097.

Related Documentation

- Junos OS MPLS for J-EX Series Switches Overview on page 3057
- Understanding Junos OS MPLS Components for J-EX Series Switches on page 3059
- Example: Configuring MPLS on J-EX Series Switches on page 3071
- Configuring MPLS on Provider Edge Switches (CLI Procedure)
- *Junos OS MPLS Applications Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

Understanding Using CoS with MPLS Networks on J-EX Series Switches

You can use class of service (CoS) within MPLS networks to prioritize certain types of traffic during periods of congestion.

J-EX Series Switches support Differentiated Service Code Point (DSCP) or IP precedence and IEEE 802.1p CoS classifiers on the customer-edge interfaces of the ingress provider edge (PE) switch. DSCP or IP precedence classifiers are used for Layer 3 packets. IEEE 802.1p is used for Layer 2 packets.

When a packet enters a customer-edge interface of the ingress PE switch, the switch associates the packet with a particular CoS servicing level prior to putting the packet onto the label-switched path (LSP). The switches within the LSP utilize the CoS value set at the ingress PE switch. The CoS value that was embedded in the DSCP, IP precedence, or IEEE 802.1 p classifier is translated and encoded in the MPLS header by means of the EXP or experimental bits.

J-EX Series switches enable a default EXP classifier and a default EXP rewrite rule. You can configure a custom EXP classifier and a custom EXP rewrite rule if you prefer. However, the switch supports only one type of EXP classifier (default or custom) and only one EXP rewrite rule (default or custom).

You do not bind the EXP classifier or the EXP rewrite rule to individual interfaces. The switch automatically and implicitly applies the default or the custom EXP classifier and the default or the custom EXP rewrite rule to the appropriate MPLS-enabled interfaces. Because rewrite rules affect only egress interfaces, the switch applies the EXP rewrite rule only to those MPLS interfaces that are transmitting MPLS packets (not to the MPLS interfaces that are receiving the packets).

This topic includes:

Guidelines for Using CoS Classifiers on CCCs

When you are configuring CoS for MPLS over circuit cross-connect (CCC), there are some additional guidelines, as follows:

- You *must* explicitly bind a CoS classifier to the CCC interface on the ingress PE switch.
- You *cannot* use more than one type of DSCP/IP precedence and not more than one type of IEEE 802.1p classifier on the CCC interfaces. Thus, if you configure one CCC interface to use DSCP1, you cannot configure another CCC interface to use DSCP2. Likewise, if you configure one CCC interface to use IEEE1, you cannot configure another

CCC interface on the same switch to use IEEE2. All the CCC interfaces on the switch must use the same DSCP classifier and the same type of IEEE 802.1p classifier.

- You *cannot* configure one CCC interface as DSCP and another CCC interface as IP precedence, because these classifier types overlap.
- You *can* configure one CCC interface as DSCP and another CCC interface as IEEE 802.1p.
- You *can* configure one CCC interface as both DSCP and IEEE 802.1p. If you configure a CCC interface with both these classifiers, the DSCP classifier is used for routing Layer 3 packets and the IEEE 802.1p classifier is used for routing Layer 2 packets.

NOTE: You can define multiple types of DSCP, IP precedence, and IEEE 802.1p on the switch and use the different classifier types for the non-CCC interfaces on the switch.

Using CoS Classifiers with IP over MPLS

When you are configuring CoS for IP over MPLS, the customer-edge interface uses the CoS configuration that has been set up for the switch as the default. You do not have to bind a classifier to the customer-edge interface in this case. There are no restrictions regarding using multiple types of DSCP, IP precedence, and IEEE 802.1p on the same switch.

- You can modify the CoS classifier for a particular interface, but it is not required.
- You can configure one interface as DSCP1 and another as DSCP2 and another and IP precedence, and so forth.

Default Classifiers and Default Rewrite Rules

The default classifiers support only two forwarding classes, **best-effort** and **network-control**, and use only two queues, **0** and **7**. However, J-EX Series switches support up to sixteen forwarding classes and eight queues. To use the additional forwarding classes and queues, create a custom classifier. To modify the code point and loss priority for a specific forwarding class, configure a rewrite rule on the switch. The default rewrite rule for EXP is enabled in the default configuration. However, the default rewrite rules for the other classifiers are not enabled in the default configuration. You can display the default classifier mappings and default rewrite mappings by entering the **show class-of-service** command on the switch.

EXP Rewrite Rules

When traffic passes from the customer-edge interface to an MPLS interface, the DSCP, IP precedence, or IEEE 802.1p CoS classifier is translated into the EXP bits within the MPLS header. You cannot disable the default EXP rewrite rule, but you can configure your own custom EXP classifier and a custom EXP rewrite rule. You cannot bind the EXP classifier to individual MPLS interfaces; the switch applies it globally to all the MPLS-enabled interfaces on the switch.

Only one EXP rewrite rule (either default or custom) is supported on a switch. The switch applies it to all the MPLS-enabled egress interfaces.

Policer

Policing helps to ensure that the amount of traffic forwarded through an LSP never exceeds the requested bandwidth allocation. During periods of congestion (when the total rate of queuing packets exceeds the rate of transmission), any new packets being sent to an interface can be dropped because there is no place to store them. You should configure a policer on the ingress PE switch:

- If you are using MPLS with CCC, you bind the policer to the LSP. You cannot bind a policer to a CCC interface.
- If you are using IP over MPLS, you bind the policer to the **inet-family** customer-edge interface. You cannot bind a policer to the LSP when you are using IP over MPLS.

Schedulers

The schedulers for using CoS with MPLS are the same as for the other CoS configurations on J-EX Series switches. Default schedulers are provided for **best-effort** and **network-control** forwarding classes. If you are using **assured-forwarding**, **expedited-forwarding**, or other custom forwarding classes, we recommend that you configure a scheduler to support that forwarding class. See “Understanding CoS Schedulers” on page 2873.

Related Documentation

- Junos OS MPLS for J-EX Series Switches Overview on page 3057
- Understanding CoS Classifiers on page 2867
- Understanding CoS Schedulers on page 2873
- Example: Configuring CoS on J-EX Series Switches on page 2883
- Configuring CoS on MPLS Provider Edge Switch Using Circuit Cross-Connect (CLI Procedure) on page 2932
- Configuring Rewrite Rules for EXP Classifiers on MPLS Networks (CLI Procedure)
- Configuring CoS on Provider Switches of an MPLS Network (CLI Procedure) on page 3106
- Defining CoS Rewrite Rules (CLI Procedure) on page 2925
- Configuring Policers to Control Traffic Rates (CLI Procedure) on page 2788

Understanding MPLS Label Operations on J-EX Series Switches

In the traditional packet-forwarding paradigm, as a packet travels from one switch to the next, an independent forwarding decision is made at each hop. The IP network header is analyzed and the next hop is chosen based on this analysis and on the information in the routing table. In an MPLS environment, the analysis of the packet header is made only once, when a packet enters the MPLS tunnel (that is, the path used for MPLS traffic).

When an IP packet enters a label-switched path (LSP), the ingress provider edge (PE) switch examines the packet and assigns it a label based on its destination, placing the label in the packet's header. The label transforms the packet from one that is forwarded based on its IP routing information to one that is forwarded based on information associated with the label. The packet is then forwarded to the next provider switch in the LSP. This switch and all subsequent switches in the LSP do not examine any of the IP routing information in the labeled packet. Rather, they use the label to look up information in their label forwarding table. They then replace the old label with a new label and forward the packet to the next switch in the path. When the packet reaches the egress PE switch, the label is removed, and the packet again becomes a native IP packet and is again forwarded based on its IP routing information.

MPLS Label Switched Paths and MPLS Labels on J-EX Series Switches

When a packet enters the MPLS network, it is assigned to an LSP. Each LSP is identified by a label, which is a short (20-bit), fixed-length value at the front of the packet. Labels are used as lookup indexes for the label forwarding table. For each label, this table stores forwarding information. Because no additional parsing or lookup is done on the encapsulated packet, MPLS supports the transmission of any other protocols within the packet payload.

NOTE: MPLS for J-EX Series switches supports only single-label packets.

Figure 81 on page 3068 shows the encoding of a single label. The encoding appears after data link layer headers, but before any network layer header.

Figure 81: Label Encoding

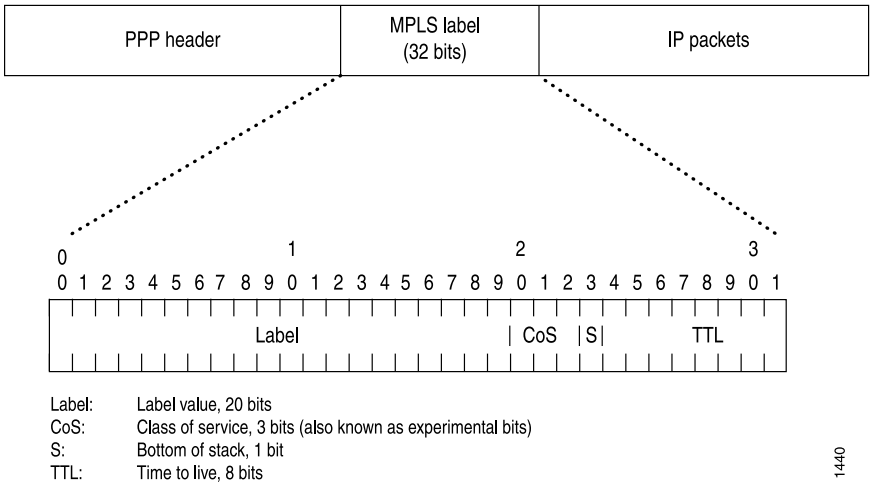

Reserved Labels

Labels range from 0 through 1,048,575. Labels 0 through 999,999 are for internal use.

Some of the reserved labels (in the 0 through 15 range) have well-defined meanings. The following reserved labels are used by J-EX Series switches:

- 0, IPv4 Explicit Null label—This value is legal only when it is the sole label entry (no label stacking). It indicates that the label must be popped on receipt. Forwarding continues based on the IP version 4 (IPv4) packet.
- 1, Router Alert label—When a packet is received with a top label value of 1, it is delivered to the local software module for processing.
- 2, IPv6 Explicit Null label—This value is legal only when it is the sole label entry (no label stacking). It indicates that the label must be popped on receipt.
- 3, Implicit Null label—This label is used in the control protocol (RSVP) only to request label popping by the downstream switch. It never actually appears in the encapsulation. Labels with a value of 3 must not be used in the data packet as real labels. No payload type (IPv4 or IPv6) is implied with this label.

MPLS Label Operations on J-EX Series Switches

J-EX Series switches support the following label operations:

- Push
- Pop
- Swap

The push operation affixes a new label to the top of the IP packet. For IPv4 packets, the new label is the first label. The time to live (TTL) field value in the packet header is derived from the IP packet header. The push operation cannot be applied to a packet that already has an MPLS label.

The pop operation removes a label from the beginning of the packet. Once the label is removed, the TTL is copied from the label into the IP packet header, and the underlying IP packet is forwarded as a native IP packet

The swap operation removes an existing MPLS label from an IP packet and replaces it with a new MPLS label, based on the following:

- Incoming interface
- Label
- Label forwarding table

Figure 82 on page 3069 shows an IP packet without a label arriving on the customer-edge interface (**ge-0/0/1**) of the ingress PE switch. The ingress PE switch examines the packet and identifies that packet's destination is the egress PE switch. The ingress PE switch applies label 100 to the packet and sends the MPLS packet to its outgoing MPLS core interface (**ge-0/0/5**). The MPLS packet is transmitted on the MPLS tunnel through the provider switch, where it arrives at interface **ge-0/0/5** with label 100. The provider switch swaps label 100 to label 200 and forwards the MPLS packet through its core interface (**ge-0/0/7**) to the next hop on the tunnel, which is the egress PE switch. The egress PE switch receives the MPLS packet through its core interface (**ge-0/0/7**), removes the MPLS label and sends the IP packet out of its customer-edge interface (**ge-0/0/1**) to a destination that is beyond the scope of the tunnel.

Figure 82: MPLS Label Swapping

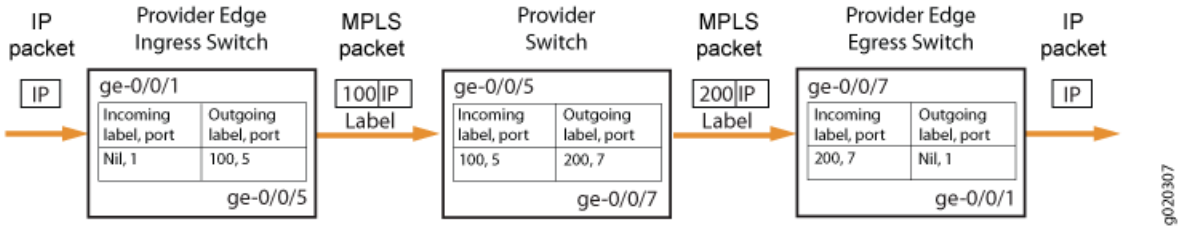

Figure 82 on page 3069 shows the path of a packet as it passes in one direction from the ingress PE switch to the egress PE switch. However, the MPLS configuration also allows traffic to travel in the reverse direction. Thus, each provider edge switch operates as both an ingress switch and an egress switch.

Ultimate and Penultimate Hop Popping

J-EX Series switches support ultimate and penultimate hop popping (that is, popping off the MPLS label) as follows:

- With circuit cross-connect (CCC), *ultimate hop popping* (UHP) is enabled by default and label 0 (IPv4 Explicit Null Label) is advertised. With UHP, the egress PE switch is responsible for popping the MPLS label at the termination of the CCC.
- With IP and MPLS, *penultimate hop popping* (PHP) is enabled by default. With PHP, the penultimate provider switch is responsible for popping the MPLS label and forwarding the traffic to the egress PE switch. The egress PE switch then performs an IP route lookup and forwards the traffic. This reduces the processing load on the egress PE switch, because it is not responsible for popping off the MPLS label.

Related Documentation

- Understanding Junos OS MPLS Components for J-EX Series Switches on page 3059
- Example: Configuring MPLS on J-EX Series Switches on page 3071
- Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect (CLI Procedure) on page 3111
- Configuring MPLS on Provider Edge Switches Using IP over MPLS (CLI Procedure) on page 3107
- Configuring MPLS on Provider Switches (CLI Procedure) on page 3102
- *Junos OS MPLS Applications Configuration Guide* at http://www.juniper.net/techpubs/software/junos/
- *Junos OS VPNs Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

CHAPTER 121

Example of MPLS Configuration

Example: Configuring MPLS on J-EX Series Switches

You can configure MPLS on J-EX Series switches to increase transport efficiency in your network. MPLS services can be used to connect various sites to a backbone network and to ensure better performance for low-latency applications such as VoIP and other business-critical functions.

To implement MPLS on J-EX Series switches, you must configure two provider edge (PE) switches—an ingress PE switch and an egress PE switch—and at least one provider (transit) switch. You can configure the customer-edge interfaces on the PE switches of the MPLS network as either circuit cross-connect (CCC) or IP (**family inet**) interfaces. This example shows how to configure an MPLS tunnel using a CCC. For information on configuring MPLS with an IP interface, see "Configuring MPLS on Provider Edge Switches Using MPLS Over IP (CLI Procedure)" on page 3107.

Requirements

This example uses the following hardware and software components:

- Three J-EX Series switches

Before you begin configuring MPLS, ensure that you have configured the routing protocol (OSPF or IS-IS) on core interface and the loopback interface on all the switches. This example includes the configuration of OSPF on the switches.

Overview and Topology

This example includes an ingress or local PE switch, an egress or remote PE switch, and one provider switch. It includes CCCs that tie the customer-edge interface of the local PE switch (PE-1) to the customer-edge interface of the remote PE switch (PE-2). It also describes how to configure the core interfaces of the PE switches and the provider switch to support the transmission of the MPLS packets. In this example, the core interfaces that connect the local PE switch and the provider switch are individual interfaces while the core interfaces that connect the remote PE switch and the provider switch are aggregated Ethernet interfaces.

NOTE:

- Core interfaces cannot be tagged VLAN interfaces.
- Core interfaces can be aggregated ethernet interfaces. This example includes a LAG between the provider switch and the remote PE switch because this type of configuration is another option you can implement. For information on configuring LAGs, see “Configuring Aggregated Ethernet Interfaces (CLI Procedure)” on page 922.

Figure 83 on page 3072 shows the topology used in this example.

Figure 83: Configuring MPLS on J-EX Series Switches

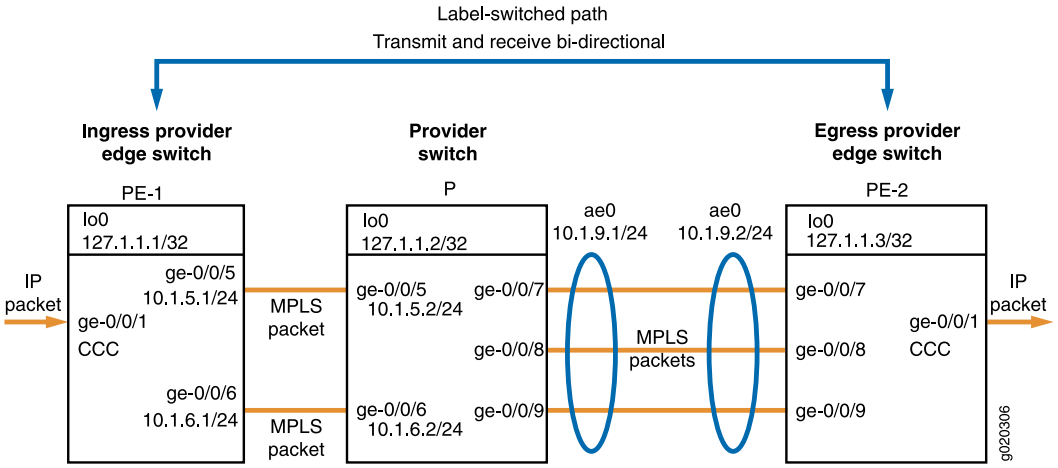

Table 406 on page 3073 shows the MPLS configuration components used for the ingress PE switch in this example.

Table 406: Components of the Ingress PE Switch in Topology for MPLS with Interface-Based CCC

Property	Settings	Description
Local PE switch hardware	J-EX Series switch	PE-1
Loopback address	**lo0 127.1.1.1/32**	Identifies PE-1 for interswitch communications.
Routing protocol	**ospf traffic-engineering**	Indicates that this switch is using OSPF as the routing protocol and that traffic engineering is enabled.
MPLS protocol and definition of label-switched path	**mpls** **label-switched-path lsp_to_pe2_ge1** **to 127.1.13**	Indicates that this PE switch is using the MPLS protocol with the specified LSP to reach the other PE switch (specified by the loopback address). The statement must also specify the core interfaces to be used for MPLS traffic.
RSVP protocol	**rsvp**	Indicates that this switch is using the RSVP protocol. The statement must specify the loopback address and the core interfaces that will be used for the RSVP session.
Interface family	**family inet** **family mpls** **family ccc**	The logical units of the core interfaces are configured to belong to both **family inet** and **family mpls**. The logical unit of the customer-edge interface is configured to belong to **family ccc**.
Customer-edge interface	**ge-0/0/1**	Interface that connects this network to devices outside the network.
Core interfaces	**ge-0/0/5.0** and **ge-0/0/6.0** with IP addresses **10.1.5.1/24** and **10.1.6.1/24**	Interfaces that connect to other switches within the MPLS network.
CCC definition	**connections** **remote-interface-switch ge-1-to-pe2** **interface ge-0/0/1.0** **transmit-lsp lsp_to_pe2_ge1** **receive-lsp lsp_to_pe1_ge1**	Associates the circuit cross-connect (CCC), **ge-0/0/1**, with the LSPs that have been defined on the local and remote PE switches.

Table 407 on page 3074 shows the MPLS configuration components used for the egress PE switch in this example.

Table 407: Components of the Egress PE Switch in Topology for MPLS with Interface-Based CCC

Property	Settings	Description
Remote PE switch hardware	J-EX Series switch	PE-2
Loopback address	**lo0 127.1.1.3/32**	Identifies PE-2 for interswitch communications.
Routing protocol	**ospf traffic-engineering**	Indicates that this switch is using OSPF as the routing protocol and that traffic engineering is enabled.
MPLS protocol and definition of label-switched path	**mpls** **label-switched-path lsp_to_pe1_ge1** **to 127.1.1.1**	Indicates that this PE switch is using the MPLS protocol with the specified label-switched path (LSP) to reach the other PE switch. The statement must also specify the core interfaces to be used for MPLS traffic.
RSVP protocol	**rsvp**	Indicates that this switch is using the RSVP protocol. The statement must specify the loopback address and the core interfaces that will be used for the RSVP session.
Interface family	**family inet** **family mpls** **family ccc**	The logical unit of the core interface is configured to belong to both **family inet** and **family mpls**. The logical unit of the customer-edge interface is configured to belong to **family ccc**.
Customer-edge interface	**ge-0/0/1**	Interface that connects this network to devices outside the network.
Core interface	**ae0** with IP address **10.1.9.2/24**	Aggregated Ethernet interface on PE-2 that connects to aggregated Ethernet interface **ae0** of the provider switch and belongs to **family mpls**.
CCC definition	**connections remote-interface-switch ge-1-to-pe1** **interface ge-0/0/1.0** **transmit-lsp lsp_to_pe1_ge1;** **receive-lsp lsp_to_pe2_ge1;**	Associates the CCC, **ge-0/0/1**, with the LSPs that have been defined on the local and remote PE switches.

Table 408 on page 3075 shows the MPLS configuration components used for the provider switch in this example.

Table 408: Components of the Provider Switch in Topology for MPLS with Interface-Based CCC

Property	Settings	Description
Provider switch hardware	J-EX Series switch	Transit switch within the MPLS network configuration.
Loopback address	**lo0 127.1.1.2/32**	Identifies provider switch for interswitch communications.
Routing protocol	**ospf traffic-engineering**	Indicates that this switch is using OSPF as the routing protocol and that traffic engineering is enabled.
MPLS protocol	**mpls**	Indicates that this switch is using the MPLS protocol. The statement must specify the core interfaces that will be used for MPLS traffic.
RSVP protocol	**rsvp**	Indicates that this switch is using the RSVP protocol. The statement must specify the loopback and the core interfaces that will be used for the RSVP session.
Interface family	**family inet** **family mpls**	The logical units for the loopback interface and core interfaces belong to **family inet**. The logical units of the core interfaces are also configured to belong to **family mpls**.
Core interfaces	**ge-0/0/5.0** and **ge-0/0/6.0** with IP addresses **10.1.5.1/24** and **10.1.6.1/24** and **ae0** with IP address **10.1.9.1/24**	Interfaces that connect the provider switch (P) to PE-1. Aggregated Ethernet interface on P that connects to aggregated Ethernet interface **ae0** of PE-2.

Configuring the Local PE Switch

CLI Quick Configuration

To quickly configure the local ingress PE switch, copy the following commands and paste them into the switch terminal window of PE-1:

```
[edit]
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/5.0
set protocols ospf area 0.0.0.0 interface ge-0/0/6.0
set protocols mpls label-switched-path lsp_to_pe2_ge1 to 127.1.1.3
```

```
set protocols mpls interface ge-0/0/5.0
set protocols mpls interface ge-0/0/6.0
set protocols rsvp interface lo0.0
set protocols rsvp interface ge-0/0/5.0
set protocols rsvp interface ge-0/0/6.0
set interfaces lo0 unit 0 family inet address 127.1.1.1/32
set interfaces ge-0/0/5 unit 0 family inet address 10.1.5.1/24
set interfaces ge-0/0/6 unit 0 family inet address 10.1.6.1/24
set interfaces ge-0/0/5 unit 0 family mpls
set interfaces ge-0/0/6 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family ccc
set protocols connections remote-interface-switch ge-1–to-pe2 interface ge-0/0/1.0
set protocols connections remote-interface-switch ge-1–to-pe2 transmit-lsp lsp_to_pe2_ge1
set protocols connections remote-interface-switch ge-1–to-pe2 receive-lsp lsp_to_pe1_ge1
```

Step-by-Step Procedure

To configure the local ingress PE switch:

1. Configure OSPF with traffic engineering enabled:

   ```
   [edit protocols]
   user@switchPE-1# set ospf traffic-engineering
   ```

2. Configure OSPF on the loopback address and core interfaces:

   ```
   [edit protocols]
   user@switchPE-1# set ospf area 0.0.0.0 interface lo0.0
   user@switchPE-1# set ospf area 0.0.0.0 interface ge-0/0/5.0
   user@switchPE-1# set ospf area 0.0.0.0 interface ge-0/0/6.0
   ```

3. Configure MPLS on the local switch with a label-switched path to the remote egress PE switch:

   ```
   [edit protocols]
   user@switchPE-1# set mpls label-switched-path lsp_to_pe2_ge1 to 127.1.1.3
   ```

4. Configure MPLS on the core interfaces:

   ```
   [edit protocols]
   user@switchPE-1# set mpls interface ge-0/0/5.0
   user@switchPE-1# set mpls interface ge-0/0/6.0
   ```

5. Configure RSVP on the loopback interface and core interfaces:

   ```
   [edit protocols]
   user@switchPE-1# set rsvp interface lo0.0
   user@switchPE-1# set rsvp interface ge-0/0/5.0
   user@switchPE-1# set rsvp interface ge-0/0/6.0
   ```

6. Configure IP addresses for the loopback interface and core interfaces:

   ```
   [edit]
   user@switchPE-1# set interfaces lo0 unit 0 family inet address 127.1.1.1/32
   user@switchPE-1# set interfaces ge-0/0/5 unit 0 family inet address 10.1.5.1/24
   user@switchPE-1# set interfaces ge-0/0/6 unit 0 family inet address 10.1.6.1/24
   ```

7. Configure **family mpls** on the logical unit of the core interface addresses:

   ```
   [edit]
   user@switchPE-1# set interfaces ge-0/0/5 unit 0 family mpls
   user@switchPE-1# set interfaces ge-0/0/6 unit 0 family mpls
   ```

8. Configure the logical unit of the customer-edge interface as a CCC:

```
[edit interfaces ge-0/0/1 unit 0]
user@PE-1# set family ccc
```

9. Configure the interface-based CCC from PE-1 to PE-2:

NOTE: You can also configure a tagged VLAN interface as a CCC. See Configuring MPLS on Provider Edge Switches (CLI Procedure).

```
[edit protocols]
user@PE-1# set connections remote-interface-switch ge-1—to-pe2 interface ge-0/0/1.0
user@PE-1# set connections remote-interface-switch ge-1—to-pe2 transmit-lsp
lsp_to_pe2_ge1
user@PE-1# set connections remote-interface-switch ge-1—to-pe2 receive-lsp
lsp_to_pe1_ge1
```

Results Display the results of the configuration:

```
user@switchPE-1> show configuration
```

```
interfaces {
    ge-0/0/1 {
        unit 0 {
            family ccc;
        }
    }
    ge-0/0/5 {
        unit 0 {
            family inet {
                address 10.1.5.1/24;
            }
            family mpls;
        }
    }
    ge-0/0/6 {
        unit 0 {
            family inet {
                address 10.1.6.1/24;
            }
            family mpls;
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 127.1.1.1/32;
            }
        }
    }
}
protocols {
    rsvp {
        interface lo0.0;
        interface ge-0/0/5.0;
        interface ge-0/0/6.0;
    }
    mpls {
```

```
    label-switched-path lsp_to_pe2_ge1 {
      to 127.1.1.3;
    }
    interface ge-0/0/5.0;
    interface ge-0/0/6.0;
  }
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
      interface lo0.0;
      interface ge-0/0/5.0;
      interface ge-0/0/6.0;
    }
  }
  connections {
    remote-interface-switch ge-1-to-pe2 {
      interface ge-0/0/1.0;
      transmit-lsp lsp_to_pe2_ge1;
      receive-lsp lsp_to_pe1_ge1;
    }
  }
```

Configuring the Remote PE Switch

CLI Quick Configuration

To quickly configure the remote PE switch, copy the following commands and paste them into the switch terminal window of PE-2:

```
[edit]
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ae0
set protocols mpls label-switched-path lsp_to_pe1_ge1 to 127.1.1.1
set protocols mpls interface ae0
set protocols rsvp interface lo0.0
set protocols rsvp interface ae0
set interfaces lo0 unit 0 family inet address 127.1.1.3/32
set interfaces ae0 unit 0 family inet address 10.1.9.2/24
set interfaces ae0 unit 0 family mpls
set interfaces ge-0/0/1 unit 0 family ccc
set protocols connections remote-interface-switch ge-1-to-pe1 interface ge-0/0/1.0
set protocols connections remote-interface-switch ge-1-to-pe1 transmit-lsp lsp_to_pe1_ge1
set protocols connections remote-interface-switch ge-1-to-pe1 receive-lsp lsp_to_pe2_ge1
```

Step-by-Step Procedure

To configure the remote PE switch (PE-2):

1. Configure OSPF with traffic engineering enabled:

   ```
   [edit protocols]
   user@switchPE-2# set ospf traffic-engineering
   ```

2. Configure OSPF on the loopback interface and core interface:

   ```
   [edit protocols]
   user@switchPE-2# set ospf area 0.0.0.0 interface lo0.0
   user@switchPE-2# set ospf area 0.0.0.0 interface ae0
   ```

3. Configure MPLS on the switch with a label-switched path to the other PE switch:

   ```
   [edit protocols]
   ```

```
user@switchPE-2# set mpls label-switched-path lsp_to_pe1_ge1 to 127.1.1.1
```

4. Configure MPLS on the core interface:

```
[edit protocols]
user@switchPE-2# set mpls interface ae0
```

5. Configure RSVP on the loopback interface and core interface:

```
[edit protocols]
ser@switchPE-2# set rsvp interface lo0.0
user@switchPE-2# set rsvp interface ae0
```

6. Configure IP addresses for the loopback interface and core interface:

```
[edit]
user@switchPE-2# set interfaces lo0 unit 0 family inet address 127.1.1.3/32
user@switchPE-2# set interfaces ae0 unit 0 family inet address 10.1.9.2/24
```

7. Configure **family mpls** on the logical unit of the core interface:

```
[edit]
user@switchPE-2# set interfaces ae0 unit 0 family mpls
```

8. Configure the logical unit of the customer-edge interface as a CCC:

```
[edit interfaces ge-0/0/1 unit 0]
user@PE-2# set family ccc
```

9. Configure the interface-based CCC from PE-2 to PE-1:

```
[edit protocols]
user@PE-2# set connections remote-interface-switch ge-1—to-pe2 interface ge-0/0/1.0
user@PE-2# set connections remote-interface-switch ge-1—to-pe2 transmit-lsp
lsp_to_pe1_ge1
user@PE-2# set connections remote-interface-switch ge-1—to-pe2 receive-lsp
lsp_to_pe2_ge1
```

Results Display the results of the configuration:

```
user@switchPE-2> show configuration

interfaces {
    ge-0/0/1 {
        unit 0 {
            family ccc;
        }
    }
    ae0 {
        unit 0 {
            family inet {
                address 10.1.9.2/24;
            }
            family mpls;
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 127.1.1.3/32;
            }
```

```
    }
  }
}
protocols {
  rsvp {
    interface lo0.0;
    interface ae0.0;
  }
  mpls {
    label-switched-path lsp_to_pe1_ge1 {
      to 127.1.1.1;
    }
    interface ae0.0;
  }
  ospf {
    traffic-engineering;
    area 0.0.0.0 {
      interface ae0.0;
    }
  }
  connections {
    remote-interface-switch ge-1-to-pe1 {
      interface ge-0/0/1.0;
      transmit-lsp lsp_to_pe1_ge1;
      receive-lsp lsp_to_pe2_ge1;
    }
  }
}
```

Configuring the Provider Switch

CLI Quick Configuration

To quickly configure the provider switch, copy the following commands and paste them into the switch terminal window:

```
[edit]
set protocols ospf traffic-engineering
set protocols ospf area 0.0.0.0 interface lo0.0
set protocols ospf area 0.0.0.0 interface ge-0/0/5.0
set protocols ospf area 0.0.0.0 interface ge-0/0/6.0
set protocols ospf area 0.0.0.0 interface ae0
set protocols mpls interface ge-0/0/5.0
set protocols mpls interface ge-0/0/6.0
set protocols mpls interface ae0
set protocols rsvp interface lo0.0
set protocols rsvp interface ge-0/0/5.0
set protocols rsvp interface ge-0/0/6.0
set protocols rsvp interface ae0
set interfaces lo0 unit 0 family inet address 127.1.1.2/32
set interfaces ge-0/0/5 unit 0 family inet address 10.1.5.1/24
set interfaces ge-0/0/6 unit 0 family inet address 10.1.6.1/24
set interfaces ae0 unit 0 family inet address 10.1.9.1/24
set interfaces ge-0/0/5 unit 0 family mpls
set interfaces ge-0/0/6 unit 0 family mpls
set interfaces ae0 unit 0 family mpls
```

Step-by-Step Procedure

To configure the provider switch:

1. Configure OSPF with traffic engineering enabled:

```
[edit protocols]
user@switchP# set ospf traffic-engineering
```

2. Configure OSPF on the loopback interface and core interfaces:

```
[edit protocols]
user@switchP# set ospf area interface lo0.0
user@switchP# set ospf area interface ge-0/0/5
user@switchP# set ospf area interface ge-0/0/6
user@switchP# set ospf area interface ae0
```

3. Configure MPLS on the core interfaces on the switch:

```
[edit protocols]
user@switchP# set mpls interface ge-0/0/5
user@switchP# set mpls interface ge-0/0/6
user@switchP# set mpls interface ae0
```

4. Configure RSVP on the loopback interface and core interfaces:

```
[edit protocols]
user@switchP# set rsvp interface lo0.0
user@switchP# set rsvp interface ge-0/0/5
user@switchP# set rsvp interface ge-0/0/6
user@switchP# set rsvp interface ae0
```

5. Configure IP addresses for the loopback and core interfaces:

```
[edit]
user@switchP# set interfaces lo0 unit 0 family inet address 127.1.1.2/32
user@switchP# set interfaces ge-0/0/5 unit 0 family inet address 10.1.5.1/24
user@switchP# set interfaces ge-0/0/6 unit 0 family inet address 10.1.6.1/24
user@switchP# set interfaces ae0 unit 0 family inet address 10.1.9.1/24
```

6. Configure **family mpls** on the logical unit of the core interface addresses:

```
[edit]
user@switchP# set interfaces ge-0/0/5 unit 0 family mpls
user@switchP# set interfaces ge-0/0/6 unit 0 family mpls
user@switchP# set interfaces ae0 unit 0 family mpls
```

Results

Display the results of the configuration:

```
user@switchP> show configuration

interfaces {
  ge-0/0/5 {
    unit 0 {
      family inet {
        address 10.1.5.1/24;
      }
      family mpls;
      }
    }
  ge-0/0/6 {
    unit 0 {
      family inet {
```

```
            address 10.1.6.1/24;
          }
          family mpls;
          }
        }
      }
      ae0 {
        unit 0 {
          family inet {
            address 10.1.9.1/24;
          }
          family mpls;
        }
      }
      lo0 {
        unit 0 {
          family inet {
            address 127.1.1.2/32;
          }
        }
      }
    protocols {
      rsvp {
        interface lo0.0;
        interface ge-0/0/5.0;
        interface ge-0/0/6.0;
        interface ae0.0;
      }
      mpls {
        interface ge-0/0/5.0;
        interface ge-0/0/6.0;
        interface ae0.0;
      }
      ospf {
        traffic-engineering;
        area 0.0.0.0 {
          interface lo0.0;
          interface ge-0/0/5.0;
          interface ge-0/0/6.0;
          interface ae0.0;
        }
      }
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying the Physical Layer on the Switches on page 3083
- Verifying the Routing Protocol on page 3083
- Verifying the Core Interfaces Being Used for the MPLS Traffic on page 3083
- Verifying RSVP on page 3084
- Verifying the Assignment of Interfaces for MPLS Label Operations on page 3084
- Verifying the Status of the CCC on page 3084

Verifying the Physical Layer on the Switches

Purpose Verify that the interfaces are up. Perform this verification task on each of the switches.

Action

```
user@switchPE-1>  interface-name terse

Interface               Admin Link Proto    Local                 Remote
ge-0/0/0                up    up
ge-0/0/0.0              up    up
ge-0/0/1.0              up    up   ccc
ge-0/0/2.0              up    up   eth-switch
ge-0/0/3.0              up    up   eth-switch
ge-0/0/4.0              up    up   eth-switch
ge-0/0/5.0              up    up   inet      10.1.5.1/24
            mpls
ge-0/0/6.0              up    up   inet      10.1.6.1/24
            mpls
```

Meaning The **show interfaces terse** command displays status information about the Gigabit Ethernet interfaces on the switch. This output verifies that the interfaces are **up**. The output for the protocol family (**Proto** column) shows that interface **ge-0/0/1.0** is configured as a circuit cross-connect. The output for the protocol family of the core interfaces (**ge-0/0/5.0** and **ge-0/0/6.0**), shows that these interfaces are configured as both **inet** and **mpls**. The **Local** column for the core interfaces shows the IP address configured for these interfaces.

Verifying the Routing Protocol

Purpose Verify the state of the configured routing protocol. Perform this verification task on each of the switches. The state must be **Full**.

Action

```
user@switchPE-1> show ospf neighbor

Address          Interface              State     ID               Pri  Dead
127.1.1.2          ge–0/0/5               Full     10.10.10.10      128    39
```

Meaning The **show ospf neighbor** command displays the status of the routing protocol. This output shows that the state is **Full**, meaning that the routing protocol is operating correctly—that is, hello packets are being exchanged between directly connected neighbors.

Verifying the Core Interfaces Being Used for the MPLS Traffic

Purpose Verify that the state of the MPLS interface is **Up**. Perform this verification task on each of the switches.

Action

```
user@switchPE-1> show mpls interface

Interface        State       Administrative groups
ge–0/0/5           Up          <none>
ge–0/0/6           Up          <none>
```

Meaning The **show mpls interface** command displays the status of the core interfaces that have been configured to belong to **family mpls**. This output shows that the interface configured to belong to **family mpls** is **Up**.

Verifying RSVP

Purpose Verify the state of the RSVP session. Perform this verification task on each of the switches.

Action

```
user@switchPE-1>  show rsvp session

Ingress RSVP: 1 sessions
To              From            State   Rt Style Labelin Labelout LSPname
127.1.13        127.1.1.1         Up       0  1 FF       -   300064 lsp_to_pe2_ge1
Total 1 displayed, Up 1, Down 0

Egress RSVP: 1 sessions
To              From            State   Rt Style Labelin Labelout LSPname
127.1.1.1         127.1.1.3         Up       0  1 FF  299968        -
lsp_to_pe1_ge1
Total 1 displayed, Up 1, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Meaning This output confirms that the RSVP sessions are **Up**.

Verifying the Assignment of Interfaces for MPLS Label Operations

Purpose Verify which interface is being used as the beginning of the CCC and which interface is being used to push the MPLS packet to the next hop. Perform this task only on the PE switches.

Action

```
user@switchPE-1>  show route forwarding-table family mpls
MPLS:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    dscd    50     1
0                  user     0                    recv    49     3
1                  user     0                    recv    49     3
2                  user     0                    recv    49     3
299776             user     0                   Pop     541     2 ge-0/0/1.0
ge-0/0/1.0  (CCC)  user     0 2.0.0.1           Push 299792   540 2   ge-0/0/5.0
```

Meaning This output shows that the CCC has been set up on interface **ge-0/0/1.0**. The switch receives ingress traffic on **ge-0/0/1.0** and pushes label **299792** onto the packet, which goes out through interface **ge-0/0/5.0**. The output also shows when the switch receives an MPLS packet with label 29976, it pops the label and sends the packet out through interface **ge-0/0/1.0**

After you have checked the local PE switch, run the same command on the remote PE switch.

Verifying the Status of the CCC

Purpose Verify the status of the CCC. Perform this task only on the PE switches.

Action

```
user@switchPE-1> show connections
CCC and TCC connections [Link Monitoring On]
Legend for status (St)              Legend for connection types
UN -- uninitialized                 if-sw:  interface switching
NP -- not present                   rmt-if: remote interface switching
WE -- wrong encapsulation           lsp-sw: LSP switching
DS -- disabled                      tx-p2mp-sw: transmit P2MP switching
Dn -- down                          rx-p2mp-sw: receive P2MP switching
-> -- only outbound conn is up
 <- -- only inbound  conn is up      Legend for circuit types
Up -- operational                   intf -- interface
RmtDn -- remote CCC down            tlsp -- transmit LSP
Restart -- restarting               rlsp -- receive LSP


Connection/Circuit               Type        St      Time last up     # Up trans
ge1-to-pe2                       rmt-if      Up      Feb 17 05:00:09  1
  ge-0/0/1.0                      intf       Up
  lsp_to_pe1_ge1                    tlsp     Up
  lsp_to_pe2_ge1                    rlsp     Up
```

Meaning

The **show connections** command displays the status of the CCC connections. This output verifies that the CCC interface and its associated transmit and receive LSPs are **Up**. After you have checked the local PE switch, run the same command on the remote PE switch.

Related Documentation

- Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect (CLI Procedure) on page 3111
- Configuring MPLS on Provider Edge Switches Using IP Over MPLS (CLI Procedure) on page 3107
- Configuring MPLS on Provider Switches (CLI Procedure) on page 3102
- Junos OS MPLS for J-EX Series Switches Overview on page 3057
- For information on the interface statement for OSPF, see the *Junos OS Routing Protocols Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.

Example: Combining CoS with MPLS on J-EX Series Switches

You can use class of service (CoS) within MPLS networks to prioritize certain types of traffic during periods of congestion. The CoS value is included within the MPLS label, which is passed through the network, enabling end-to-end CoS across the network.

MPLS services are often used to ensure better performance for low-latency applications such as VoIP and other business-critical functions. These applications place specific demands on a network for successful transmission. CoS gives you the ability to control the mix of bandwidth, delay, jitter, and packet loss while taking advantage of the MPLS labeling mechanism.

This example shows how to configure CoS on an MPLS network that is using a unidirectional circuit cross-connect (CCC) from the ingress provider edge (PE) switch to the egress PE switch. for the customer-edge interface of the ingress provider edge (PE) switch. It describes adding the configuration of CoS components to the ingress PE switch,

the egress PE switch, and the core provider switches of the existing MPLS network. Because of the unidirectional configuration, the DSCP classifier needs to be configured only on the ingress PE switch.

Requirements

This example uses the following hardware and software components:

- Three J-EX Series switches

Before you configure CoS with MPLS, be sure you have:

Configured an MPLS network with two PE switches and one provider switch. See "Example: Configuring MPLS on J-EX Series Switches" on page 3071. This example assumes that an MPLS network has been configured using a cross circuit-connect (CCC).

Overview and Topology

This example describes adding custom classifiers and custom rewrite rules to switches in an MPLS network that is using MPLS over CCC.

It is a unidirectional configuration. Therefore, you need to configure custom classifiers and custom rewrite rules as follows:

- On the ingress PE switch: custom DSCP classifier and custom EXP rewrite rule
- On the egress PE switch: custom EXP classifier
- On the provider switch: customer EXP classifier and custom EXP rewrite rule

NOTE: You can also configure schedulers and shapers as needed. If you are using assured-forwarding, expedited-forwarding, or other custom forwarding classes, we recommend that you configure a scheduler to support that forwarding class. See "Defining CoS Schedulers (CLI Procedure)" on page 2920.

The example creates a custom DSCP classifier (**dscp1**) on the ingress PE switch and binds this classifier to the CCC interface. It includes configuration of a policer on the ingress PE switch. The policer is applied as a filter on the label-switched path (LSP) **lsp_to_pe2_ge1**(created in "Example: Configuring MPLS on J-EX Series Switches" on page 3071) to ensure that the amount of traffic forwarded through the LSP never exceeds the requested bandwidth allocation.

This example creates a custom EXP rewrite rule (**exp1**) on the ingress PE switch, specifying a loss-priority and code point to be used for the expedited-forwarding class as the packet travels through the LSP. The switch applies this custom rewrite rule on the core interfaces **ge-0/0/5.0** and **ge-0/0/6.0**, which are the egress interfaces for this switch.

Table 409 on page 3087 shows the CoS configuration components added to the ingress PE switch.

Table 409: CoS Configuration Components on the Ingress PE Switch

Property	Settings	Description
Local PE switch hardware	J-EX Series switch	PE-1
Policing filter configured and applied to the LSP.	**policing filter mypolicer** **filter myfilter**	Name of the rate-limiting policer. Name of the filter, which refers to the policer
Custom DSCP classifier	**dscp1**	Specifies the name of the custom DSCP classifier
Custom EXP rewrite rule	**e1**	Name of the custom EXP rewrite rule.
Customer-edge interface	**ge-0/0/1.0**	Interface that receives packets from devices outside the network. The custom DSCP classifier must be specified on this CCC interface.
Core interfaces	**ge-0/0/5.0** and **ge-0/0/6.0**	Interfaces that transmit MPLS packets to other switches within the MPLS network. The EXP rewrite rule is applied implicitly to these interfaces.

Table 410 on page 3087 shows the CoS configuration components added to the egress PE switch in this example.

Table 410: CoS Configuration Components of the Egress PE Switch

Property	Settings	Description
Remote provider edge switch hardware	J-EX Series switch	PE-2
Custom EXP classifier	**exp1**	Name of custom EXP classifier
Customer-edge interface	**ge-0/0/1.0**	Interface that transmits packets from this network to devices outside the network. No CoS classifier is specified for this interface. A scheduler can be specified.

Table 410: CoS Configuration Components of the Egress PE Switch *(continued)*

Property	Settings	Description
Core interfaces	**ge-0/0/7.0** and **ge-0/0/8.0**	Core interfaces on PE-2 that receive MPLS packets from the provider switch. The EXP classifier is enabled by default on the switch and applied implicitly to these interfaces.

Table 411 on page 3088 shows the MPLS configuration components used for the provider switch in this example.

Table 411: CoS Configuration Components of the Provider Switch

Property	Settings	Description
Provider switch hardware	J-EX Series switch	Transit switch within the MPLS network configuration.
Custom EXP classifier	**exp1**	Name of the custom EXP classifier.
Custom EXP rewrite rule	**e1**	Name of the custom EXP rewrite rule.
Core interfaces receiving packets from other MPLS switches.	**ge-0/0/5.0** and **ge-0/0/6.0**	Interfaces that connect the provider switch to the ingress PE switch (PE-1). The EXP classifier is enabled by default on the switch and applied implicitly to these interfaces.
Core interfaces transmitting packets to other switches within the MPLS network.	**ge-0/0/7.0** and **ge-0/0/8.0**	Interfaces that transmit packets to the egress PE (PE-2). The EXP rewrite rule is applied implicitly on these interfaces. Schedulers can also be specified and will be applied to these interfaces.

Configuring the Local PE Switch

CLI Quick Configuration

To quickly configure a custom DSCP classifier, custom EXP rewrite rule, and a policer on the local PE switch, copy the following commands and paste them into the switch terminal window of PE-1:

```
[edit]
set class-of-service classifiers dscp dscp1 import default
set class-of-service classifiers dscp dscp1 forwarding-class expedited-forwarding loss-priority low code-points 000111
set class-of-service rewrite-rules exp e1 forwarding-class expedited-forwarding loss-priority low code-point 111
set class-of-service interfaces ge-0/0/1 unit 0 classifier dscp1
set firewall policer mypolicer if-exceeding bandwidth-limit 500m
set firewall policer mypolicer if-exceeding burst-size-limit 33553920
set firewall policer mypolicer then discard
set firewall family any filter myfilter term t1 then policer mypolicer
set protocols mpls label-switched-path lsp_to_pe2_ge1 to 127.1.1.3 policing filter myfilter
```

Step-by-Step Procedure

To configure a custom DSCP classifier, custom EXP rewrite rule, and a policer on the ingress PE switch:

1. Import the default DSCP classifier classes to the custom DSCP classifier that you are creating:

```
[edit class-of-service]
user@switch# set classifiers dscp dscp1 import default
```

2. Add the expedited-forwarding class to this custom DSCP classifier, specifying a loss priority and code point:

```
[edit class-of-service]
user@switch# set classifiers dscp dscp1 forwarding-class expedited-forwarding
loss-priority low code-points 000111
```

3. Specify the values for the custom EXP rewrite rule, **e1**:

```
[edit class-of-service]
user@switch#  set rewrite-rules exp e1 forwarding-class expedited-forwarding
loss-priority low code-point 111
```

4. Bind the DSCP classifier to the CCC interface:

```
[edit ]
user@switch#  set class-of-service interfaces ge-0/0/1 unit 0 classifier dscp1
```

5. Specify the number of bits per second permitted, on average, for the firewall policer, which will later be applied to the LSP:

```
[edit firewall]
set policer mypolicer if-exceeding bandwidth-limit 500m
```

6. Specify the maximum size permitted for bursts of data that exceed the given bandwidth limit for this policer:

```
[edit firewall policer]
set mypolicer if-exceeding burst-size-limit 33553920
```

7. Discard traffic that exceeds the rate limits for this policer:

```
[edit firewall policer]
set mypolicer then discard
```

8. To reference the policer, configure a filter term that includes the policer action:

```
[edit firewall]
user@switch#  set family any filter myfilter term t1 then policer mypolicer
```

9. Apply the filter to the LSP:

```
[edit protocols mpls]
set label-switched-path lsp_to_pe2_ge1 policing filter myfilter
```

Results

Display the results of the configuration:

```
[edit]
 user@switch# show
class-of-service {
  classifiers {
    dscp dscp1 {
```

```
        import default;
        forwarding-class expedited-forwarding {
          loss-priority low code-points 000111;
        }
      }
    }
    interfaces {
      ge-0/0/1 {
        unit 0 {
          classifiers {
            dscp dscp1;
          }
        }
      }
    }
  rewrite-rules {
    exp e1 {
      forwarding-class expedited-forwarding {
        loss-priority low code-point 111;
      }
    }
    }
}
firewall {
  family any {
    filter myfilter {
      term t1 {
        then policer mypolicer;
      }
    }
  }
  policer mypolicer {
    if-exceeding {
      bandwidth-limit 500m;
      burst-size-limit 33553920;
    }
    then discard;
  }
}
```

Configuring the Remote PE Switch

CLI Quick Configuration

To quickly configure a custom EXP classifier on the remote PE switch, copy the following commands and paste them into the switch terminal window of PE-2:

```
[edit]
set class-of-service classifiers exp exp1 import default
set class-of-service classifiers exp exp1 forwarding-class expedited-forwarding loss-priority low
code-points 010
```

Step-by-Step Procedure

To configure a custom EXP classifier on the egress PE switch:

1. Import the default EXP classifier classes to the custom EXP classifier that you are creating:

```
[edit class-of-service]
user@switch# set classifiers exp exp1 import default
```

2. Add the expedited-forwarding class to this custom EXP classifier, specifying a loss priority and code point:

```
[edit class-of-service]
user@switch# set classifiers exp exp1 forwarding-class expedited-forwarding loss-priority low code-points 010
```

Results Display the results of the configuration:

```
[edit]
 user@switch# show
class-of-service {
  classifiers {
    exp exp1 {
      import default;
      forwarding-class expedited-forwarding {
        loss-priority low code-points 010;
      }
    }
```

Configuring the Provider Switch

CLI Quick Configuration To quickly configure a custom EXP classifier and a custom EXP rewrite rule on the provider switch, copy the following commands and paste them into the switch terminal window of the provider switch:

```
[edit]
set class-of-service classifiers exp exp1 import default
set class-of-service classifiers exp exp1 forwarding-class expedited-forwarding loss-priority low code-points 010
set class-of-service rewrite-rules exp e1 forwarding-class expedited-forwarding loss-priority low code-point 111
```

Step-by-Step Procedure To configure a custom EXP classifier and a custom EXP rewrite rule on the provider switch:

1. Import the default EXP classifier classes to the custom EXP classifier that you are creating:

```
[edit class-of-service]
user@switch#  set classifiers exp exp1 import default
```

2. Add the expedited-forwarding class to this custom EXP classifier, specifying a loss priority and code point:

```
[edit class-of-service]
user@switch# set classifiers exp exp1 forwarding-class expedited-forwarding loss-priority low code-points 010
```

3. Specify the values for the custom EXP rewrite rule, **e1**:

```
[edit class-of-service]
user@switch#  set rewrite-rules exp e1 forwarding-class expedited-forwarding loss-priority low code-point 111
```

Results Display the results of the configuration:

```
[edit]
```

```
    user@switch# show
  class-of-service {
    classifiers {
      exp exp1 {
        import default;
        forwarding-class expedited-forwarding {
          loss-priority low code-points 010;
        }
      }
    }
    rewrite-rules {
      exp e1 {
        forwarding-class expedited-forwarding {
          loss-priority low code-point 111;
        }
      }
    }
  }
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying That the Policer Firewall Filter Is Operational on page 3092
- Verifying That the CoS Classifiers Are Going to the Right Queue on page 3092
- Verifying the CoS Forwarding Table Mapping on page 3095
- Verifying the Rewrite Rules on page 3095

Verifying That the Policer Firewall Filter Is Operational

Purpose Verify the operational state of the policer that is configured on the ingress PE switch.

Action

```
user@switch> show firewall
Filter: myfilter
Policers:
Name                                                Packets
mypolicer-t1                                              0
```

Meaning This output shows that the firewall filter **mypolicer** has been created.

Verifying That the CoS Classifiers Are Going to the Right Queue

Purpose Verify that the CoS classifiers are going to the right queue.

Action

```
user@switch> show class-of-service forwarding-table classifier

Classifier table index: 7, # entries: 64, Table type: DSCP
Entry #   Code point   Forwarding-class #   PLP
   0        000000         0       0
   1        000001         0       0
   2        000010         0       0
   3        000011         0       0
   4        000100         0       0
   5        000101         0       0
   6        000110         0       0
```

```
  7         000111         0       0
  8         001000         0       0
  9         001001         0       0
 10         001010         0       0
 11         001011         0       0
 12         001100         0       0
 13         001101         0       0
 14         001110         0       0
 15         001111         0       0
 16         010000         0       0
 17         010001         0       0
 18         010010         0       0
 19         010011         0       0
 20         010100         0       0
 21         010101         0       0
 22         010110         0       0
 23         010111         0       0
 24         011000         0       0
 25         011001         0       0
 26         011010         0       0
 27         011011         0       0
 28         011100         0       0
 29         011101         0       0
 30         011110         0       0
 31         011111         0       0
 32         100000         0       0
 33         100001         0       0
 34         100010         0       0
 35         100011         0       0
 36         100100         0       0
 37         100101         0       0
 38         100110         0       0
 39         100111         0       0
 40         101000         0       0
 41         101001         0       0
 42         101010         0       0
 43         101011         0       0
 44         101100         0       0
 45         101101         0       0
 46         101110         0       0
 47         101111         0       0
 48         110000         3       0
 49         110001         3       0
 50         110010         3       0
 51         110011         3       0
 52         110100         3       0
 53         110101         3       0
 54         110110         3       0
 55         110111         3       0
 56         111000         3       0
 57         111001         3       0
 58         111010         3       0
 59         111011         3       0
 60         111100         3       0
 61         111101         3       0
 62         111110         3       0
 63         111111         3       0

Classifier table index: 11, # entries: 8, Table type: IEEE 802.1
Entry #   Code point   Forwarding-class #   PLP
   0          000        0       0
```

```
   1           001        0       0
   2           010        0       0
   3           011        0       0
   4           100        0       0
   5           101        0       0
   6           110        3       0
   7           111        3       0

Classifier table index: 12, # entries: 8, Table type: IPv4 precedence
Entry #   Code point   Forwarding-class #   PLP
   0           000        0       0
   1           001        0       0
   2           010        0       0
   3           011        0       0
   4           100        0       0
   5           101        0       0
   6           110        3       0
   7           111        3       0

Classifier table index: 16, # entries: 8, Table type: Untrust
Entry #   Code point   Forwarding-class #   PLP
   0           000        0       0
   1           001        0       0
   2           010        0       0
   3           011        0       0
   4           100        0       0
   5           101        0       0
   6           110        0       0
   7           111        0       0

Classifier table index: 9346, # entries: 64, Table type: DSCP
Entry #   Code point   Forwarding-class #   PLP
   0        000000        0       0
   1        000001        0       0
   2        000010        0       0
   3        000011        0       0
   4        000100        0       0
   5        000101        0       0
   6        000110        0       0
   7        000111        1       0
   8        001000        0       0
   9        001001        0       0
  10        001010        0       0
  11        001011        0       0
  12        001100        0       0
  13        001101        0       0
  14        001110        0       0
  15        001111        0       0
  16        010000        0       0
  17        010001        0       0
  18        010010        0       0
  19        010011        0       0
  20        010100        0       0
  21        010101        0       0
  22        010110        0       0
  23        010111        0       0
  24        011000        0       0
  25        011001        0       0
  26        011010        0       0
  27        011011        0       0
  28        011100        0       0
```

```
29          011101          0          0
30          011110          0          0
31          011111          0          0
32          100000          0          0
33          100001          0          0
34          100010          0          0
35          100011          0          0
36          100100          0          0
37          100101          0          0
38          100110          0          0
39          100111          0          0
40          101000          0          0
41          101001          0          0
42          101010          0          0
43          101011          0          0
44          101100          0          0
45          101101          0          0
46          101110          0          0
47          101111          0          0
48          110000          3          0
49          110001          3          0
50          110010          3          0
51          110011          3          0
52          110100          3          0
53          110101          3          0
54          110110          3          0
55          110111          3          0
56          111000          3          0
57          111001          3          0
58          111010          3          0
59          111011          3          0
60          111100          3          0
61          111101          3          0
62          111110          3          0
63          111111          3          0
```

Meaning This output shows that a new DSCP classifier has been created, index **9346**, on the ingress PE switch (PE-1).

Verifying the CoS Forwarding Table Mapping

Purpose For each logical interface, display either the table index of the classifier for a given code point type or the queue number (if it is a fixed classification) in the forwarding table.

Action

```
user@switch>show class-of-service forwarding-table classifier mapping

                          Table Index/
Interface      Index       Q num      Table type
ge-0/0/1.0        92        9346      DSCP
```

Meaning The results show that the new DSCP classifier, index number **9346**, is bound to interface **ge-0/0/1.0**.

Verifying the Rewrite Rules

Purpose Display mapping of the queue number and loss priority to code point value for each rewrite rule as it exists in the forwarding table.

Action

```
user@switch>show class-of-service forwarding-table rewrite-rule

Rewrite table index: 31, # entries: 4, Table type: DSCP
FC#     Low bits  State      High bits  State
0       000000  Enabled       000000  Enabled
1       101110  Enabled       101110  Enabled
2       001010  Enabled       001100  Enabled
3       110000  Enabled       111000  Enabled

Rewrite table index: 34, # entries: 4, Table type: IEEE 802.1
FC#     Low bits  State      High bits  State
0          000  Enabled          001  Enabled
1          010  Enabled          011  Enabled
2          100  Enabled          101  Enabled
3          110  Enabled          111  Enabled

Rewrite table index: 35, # entries: 4, Table type: IPv4 precedence
FC#     Low bits  State      High bits  State
0          000  Enabled          000  Enabled
1          101  Enabled          101  Enabled
2          001  Enabled          001  Enabled
3          110  Enabled          111  Enabled

Rewrite table index: 9281, # entries: 1, Table type: EXP
FC#     Low bits  State      High bits  State
1          111  Enabled          000  Disabled
```

Meaning This output shows that a new EXP classifier with the index number **9281** has been created.

Related Documentation

- Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect (CLI Procedure) on page 3111
- Configuring MPLS on Provider Edge Switches Using IP Over MPLS (CLI Procedure) on page 3107
- Understanding Using CoS with MPLS Networks on J-EX Series Switches on page 2880
- Monitoring CoS Forwarding Classes on page 2936

CHAPTER 122

Configuring MPLS

Configuring Path Protection in an MPLS Network (CLI Procedure)

The Junos OS implementation of MPLS on J-EX Series switches provides path protection as a mechanism for protecting against label switched path (LSP) failures. Path protection reduces the time required to recalculate a route in case of a failure within the MPLS tunnel. You configure path protection on the ingress provider edge switch in your MPLS network. You do not configure the egress provider edge switch or the provider switches for path protection. You can explicitly specify which provider switches are used for the primary and secondary paths, or you can let the software calculate the paths automatically.

Before you configure path protection, be sure you have:

- Configured an ingress provider edge switch and an egress provider edge switch. See Configuring MPLS on Provider Edge Switches (CLI Procedure).
- Configured at least one provider (transit) switch. See “Configuring MPLS on Provider Switches (CLI Procedure)” on page 3102.
- Verified the configuration of your MPLS network. See “Verifying That MPLS Is Working Correctly” on page 3115.

To configure path protection, complete the following tasks on the ingress provider edge switch:

1. Configuring the Primary Path on page 3099
2. Configuring the Secondary Path on page 3099
3. Configuring the Revert Timer on page 3100

Configuring the Primary Path

The **primary** statement creates the primary path, which is the LSP's preferred path. The **secondary** statement creates an alternative path if the primary path can no longer reach the egress provider edge switch.

In the tasks described in this topic, the ***lsp-name*** has already been configured on the ingress provider edge switch as **lsp_to_240** and the loopback interface address on the remote provider edge switch has already been configured as **127.0.0.8**.

When the software switches from the primary to the secondary path, it continuously attempts to revert to the primary path, switching back to it when it is again reachable but no sooner than the retry time specified in the **revert-timer** statement.

You can configure zero primary paths or one primary path. If you do not configure a primary path, the first secondary path (if a secondary path has been configured) is selected as the path. If you do not specify any named paths, or if the path that you specify is empty, the software makes all routing decisions necessary for the packets to reach the egress provider edge switch.

To configure a primary path:

1. Create the primary path for the LSP:

```
[edit protocols mpls label-switched-path lsp_to_240 to 127.0.0.8]
user@switch# set primary primary_path_lsp_to_240
```

2. Configure an explicit route for the primary path by specifying the IP address of the loopback interface or the switch IP address or hostname of each switch used in the MPLS tunnel. You can specify the link types as either **strict** or **loose** in each **path** statement. If the link type is **strict**, the LSP must go to the next address specified in the **path** statement without traversing other switches. If the link type is **loose**, the LSP can traverse through other switches before reaching this switch. This configuration uses the default **strict** designation for the paths.

NOTE: You can enable path protection without specifying which provider switches are used. If you do not list the specific provider switches to be used for the MPLS tunnel, the switch calculates the route.

TIP: Do not include the ingress provider edge switch in these statements. List the IP address of the loopback interface or switch address or hostname of all other switch hops in sequence, ending with the egress provider edge switch.

```
[edit protocols mpls label-switched-path lsp_to_240 to 127.0.0.8]
user@switch# set path primary_path_lsp_to_240 127.0.0.2
user@switch# set path primary_path_lsp_to_240 127.0.0.3
user@switch# set path primary_path_lsp_to_240 127.0.0.8
```

Configuring the Secondary Path

You can configure zero or more secondary paths. All secondary paths are equal, and the software tries them in the order that they are listed in the configuration. The software does not attempt to switch among secondary paths. If the first secondary path in the configuration is not available, the next one is tried, as so on. To create a set of equal paths, specify secondary paths without specifying a primary path. If you do not specify any named paths, or if the path that you specify is empty, the software makes all routing decisions necessary to reach the egress provider edge switch.

To configure the secondary path:

1. Create a secondary path for the LSP:

```
[edit protocols mpls label-switched-path lsp_to_240 to 127.0.0.8]
user@switch# set secondary secondary_path_lsp_to_240 standby
```

2. Configure an explicit route for the secondary path by specifying the IP address of the loopback interface or the switch IP address or hostname of each switch used in the MPLS tunnel. You can specify the link types as either **strict** or **loose** in each **path** statement. This configuration uses the default **strict** designation for the paths.

TIP: Do not include the ingress provider edge switch in these statements. List the IP address of the loopback interface or switch address or hostname of all other switch hops in sequence, ending with the egress provider edge switch.

```
[edit protocols mpls label-switched-path lsp_to_240 to 127.0.0.8]
user@switch# set path secondary_path_lsp_to_240 127.0.0.4
user@switch# set path primary_path_lsp_to_240 127.0.0.8
```

Configuring the Revert Timer

For LSPs configured with both primary and secondary paths, you can optionally configure a revert timer. If the primary path goes down and traffic is switched to the secondary path, the revert timer specifies the amount of time (in seconds) that the LSP must wait before it can revert traffic back to the primary path. If the primary path experiences any connectivity problems or stability problems during this time, the timer is restarted.

TIP: If you do not explicitly configure the revert timer, it is set by default to 60 seconds.

To configure the revert timer for LSPs configured with primary and secondary paths:

- For all LSPs on the switch:

```
[edit protocols mpls]
user@switch#  set revert-timer 120
```

- For a specific LSP on the switch:

```
[edit protocols mpls label-switched-path]
user@switch#  set lsp_to_240 revert-timer 120
```

Related Documentation

- Understanding MPLS and Path Protection on J-EX Series Switches on page 3063

Configuring MPLS on Provider Switches (CLI Procedure)

Junos OS MPLS for J-EX Series switches supports Layer 2 protocols and Layer 2 virtual private networks (VPNs). You can configure MPLS on J-EX Series switches to increase transport efficiency in your network. MPLS services can be used to connect various sites to a backbone network and to ensure better performance for low-latency applications such as VoIP and other business-critical functions.

NOTE: You can use class of service (CoS) on MPLS networks. For further information, see “Understanding Using CoS with MPLS Networks on J-EX Series Switches” on page 2880.

To implement MPLS on J-EX Series switches, you must configure at least one provider switch as a transit switch for the MPLS packets. The configuration of all the provider switches is the same regardless of whether the provider edge (PE) switches are using circuit cross-connect (CCC) or using MPLS over IP for the customer-edge interfaces.

To configure the provider switch, complete the following tasks:

1. Enable the routing protocol (OSPF or IS-IS) on the loopback interface and on the core interfaces:

   

   NOTE: You can use the switch address as an alternative to the loopback interface.

   ```
   [edit protocols]
   user@switch# set ospf area 0.0.0.0 interface lo0.0
   user@switch# set ospf area 0.0.0.0 interface ge-0/0/5.0
   user@switch# set ospf area 0.0.0.0 interface ge-0/0/6.0
   user@switch# set ospf area 0.0.0.0 interface ae0
   ```

2. Enable traffic engineering for the routing protocol (OSPF or IS-IS):

   ```
   [edit protocols]
   user@switch# set ospf traffic-engineering
   ```

3. Enable MPLS within the **protocols** stanza and apply it to the core interfaces:

   ```
   [edit protocols]
   user@switch# set mpls interface ge-0/0/5.0
   user@switch# set mpls interface ge-0/0/6.0
   user@switch# set mpls interface ae0
   ```

4. Configure RSVP on the loopback interface and the core interfaces:

   ```
   [edit protocols]
   user@switch# set rsvp interface lo0.0
   user@switch# set rsvp interface ge-0/0/5.0
   user@switch# set rsvp interface ge-0/0/6.0
   user@switch# set rsvp interface ae0
   ```

5. Configure an IP address for the loopback interface and for the core interfaces:

```
[edit]
user@switch# set interfaces lo0 unit 0 family inet address 127.1.1.1/32
user@switch# set interfaces ge-0/0/5 unit 0 family inet address 10.1.5.1/24
user@switch# set interfaces ge-0/0/6 unit 0 family inet address 10.1.6.1/24
user@switch# set interfaces ae0 unit 0 family inet address 10.1.9.2/24
```

6. Configure **family mpls** on the logical units of the core interfaces:

```
[edit]
user@switch# set interfaces ge-0/0/5 unit 0 family mpls
user@switch# set interfaces ge-0/0/6 unit 0 family mpls
user@switch# set interfaces ae0 unit 0 family mpls
```

NOTE: You can enable family mpls on either individual interfaces or aggregated Ethernet interfaces. You cannot enable it on tagged VLAN interfaces.

Related Documentation

- Example: Configuring MPLS on J-EX Series Switches on page 3071
- Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect (CLI Procedure) on page 3111
- Configuring MPLS on Provider Edge Switches Using MPLS Over IP (CLI Procedure) on page 3107
- Configuring an OSPF Network (J-Web Procedure) on page 1435
- Understanding Junos OS MPLS Components for J-EX Series Switches on page 3059

Configuring CoS on MPLS Provider Edge Switch Using IP Over MPLS (CLI Procedure)

You can use class of service (CoS) within MPLS networks to prioritize certain types of traffic during periods of congestion. This topic describes configuring CoS components on a provider edge (PE) switch that is using IP Over MPLS.

This task describes now to create a custom DSCP classifier and a custom EXP rewrite rule on the ingress PE switch. It includes configuring a policer firewall filter and applying it to the customer-edge interface of the ingress PE switch. The policer firewall filter ensures that the amount of traffic forwarded through the MPLS tunnel never exceeds the requested bandwidth allocation.

For this procedure, we assume that the switch has already been configured for MPLS. See "Configuring MPLS on Provider Edge Switches Using MPLS Over IP (CLI Procedure)" on page 3107.

1. Import the default DSCP classifier classes to the custom DSCP classifier that you are creating:

   ```
   [edit class-of-service]
   user@switch#set classifiers dscp dscp1 import default
   ```

2. Add the expedited-forwarding class to this custom DSCP classifier, specifying a loss priority and code point:

   ```
   [edit class-of-service]
   user@switch#set classifiers dscp dscp1 forwarding-class expedited-forwarding
   loss-priority low code-points 000111
   ```

3. Specify the values for the custom EXP rewrite rule, **e1**:

   ```
   [edit class-of-service]
   user@switch#  set rewrite-rules exp e1 forwarding-class expedited-forwarding
   loss-priority low code-point 111
   ```

4. Specify the number of bits per second permitted, on average, for the firewall policer, which will later be applied to the customer-edge-interface:

   ```
   [edit firewall]
   set policer mypolicer if-exceeding bandwidth-limit 500m
   ```

5. Specify the maximum size permitted for bursts of data that exceed the given bandwidth limit for this policer:

   ```
   [edit firewall policer]
   set mypolicer if-exceeding burst-size-limit 33553920
   ```

6. Discard traffic that exceeds the rate limits for this policer:

   ```
   [edit firewall policer]
   set mypolicer then discard
   ```

7. To reference the policer, configure a filter term that includes the policer action:

   ```
   [edit firewall]
   user@switch#  set family inet filter myfilter term t1 then policer mypolicer
   ```

8. Apply the filter to the customer-edge interface:

```
[edit interfaces]
user@switch# set ge-2/0/3 unit 0 family inet address 121.121.121.1/16 policing filter
myfilter
```

NOTE: You can also configure schedulers and shapers as needed. See “Defining CoS Schedulers (CLI Procedure)” on page 2920.

Related Documentation

- Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect (CLI Procedure) on page 3111
- Assigning CoS Components to Interfaces (CLI Procedure) on page 2928
- Configuring Policers to Control Traffic Rates (CLI Procedure) on page 2788
- Understanding the Use of Policers in Firewall Filters on page 2752

Configuring CoS on MPLS Provider Edge Switch Using Circuit Cross-Connect (CLI Procedure)

You can use class of service (CoS) within MPLS networks to prioritize certain types of traffic during periods of congestion. This topic describes configuring CoS components on provider edge (PE) switch that is using MPLS over circuit-cross connect (CCC).

NOTE: If you are using MPLS with CCC, you can use only one type of DSCP/IP precedence and only one type of IEEE 802.1p on the CCC interfaces.

This procedure creates a custom DSCP classifier and a custom EXP rewrite rule on the ingress PE. It also enables a policer on the label-switched path (LSP) of the ingress PE to ensure that the amount of traffic forwarded through the LSP never exceeds the requested bandwidth allocation.

1. Import the default DSCP classifier classes to the custom DSCP classifier that you are creating:

   ```
   [edit class-of-service]
   user@switch#set classifiers dscp dscp1 import default
   ```

2. Add the expedited-forwarding class to this custom DSCP classifier, specifying a loss priority and code point:

   ```
   [edit class-of-service]
   user@switch#set classifiers dscp dscp1 forwarding-class expedited-forwarding
   loss-priority low code-points 000111
   ```

3. Specify the values for the custom EXP rewrite rule, **e1**:

   ```
   [edit class-of-service]
   user@switch# set rewrite-rules exp e1 forwarding-class expedited-forwarding
   loss-priority low code-point 111
   ```

4. Bind the DSCP classifier to the CCC interface:

```
[edit ]
user@switch# set class-of-service interfaces ge-0/0/1 unit 0 classifier dscp1
```

5. Specify the number of bits per second permitted, on average, for the firewall policer, which will later be applied to the LSP:

```
[edit firewall]
set policer mypolicer if-exceeding bandwidth-limit 500m
```

6. Specify the maximum size permitted for bursts of data that exceed the given bandwidth limit for this policer:

```
[edit firewall policer]
set mypolicer if-exceeding burst-size-limit 33553920
```

7. Discard traffic that exceeds the rate limits for this policer:

```
[edit firewall policer]
set mypolicer then discard
```

8. To reference the policer, configure a filter term that includes the policer action:

```
[edit firewall]
user@switch# set family any filter myfilter term t1 then policer mypolicer
```

9. Apply the filter to the LSP:

```
[edit protocols mpls]
set label-switched-path lsp_to_pe2_ge1 policing filter myfilter
```

NOTE: You can also configure schedulers and shapers as needed. See “Defining CoS Schedulers (CLI Procedure)” on page 2920.

Related Documentation

- Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect (CLI Procedure) on page 3111
- Assigning CoS Components to Interfaces (CLI Procedure) on page 2928
- Configuring Policers to Control Traffic Rates (CLI Procedure) on page 2788
- Understanding the Use of Policers in Firewall Filters on page 2752

Configuring CoS on Provider Switches of an MPLS Network (CLI Procedure)

You can add class-of-service (CoS) components to your MPLS networks on J-EX Series switches to achieve end-to-end Differentiated Services to match your specific business requirements. The configuration of CoS components on the provider switches is the same regardless of whether the provider edge (PE) switches are using MPLS over CCC or IP over MPLS.

This task shows how to configure a custom EXP classifier and custom EXP rewrite rule on the provider switch.

1. Import the default EXP classifier classes to the custom EXP classifier that you are creating:

```
[edit class-of-service]
user@switch# set classifiers exp exp1 import default
```

2. Add the expedited-forwarding class to this custom EXP classifier, specifying a loss priority and code point:

```
[edit class-of-service]
user@switch#set classifiers exp exp1 forwarding-class expedited-forwarding loss-priority low code-points 010
```

3. Specify the values for the custom EXP rewrite rule, **e1**:

```
[edit class-of-service]
user@switch# set rewrite-rules exp e1 forwarding-class expedited-forwarding loss-priority low code-point 111
```

NOTE: You can also configure schedulers and shapers as needed. See “Defining CoS Schedulers (CLI Procedure)” on page 2920.

Related Documentation

- Example: Combining CoS with MPLS on J-EX Series Switches on page 2883

Configuring MPLS on Provider Edge Switches Using IP Over MPLS (CLI Procedure)

You can configure MPLS on J-EX Series switches to increase transport efficiency in your network. MPLS services can be used to connect various sites to a backbone network or to ensure better performance for low-latency applications such as VoIP and other business-critical functions.

To implement MPLS on J-EX Series switches, you must configure two provider edge (PE) switches—an ingress PE switch and an egress PE switch—and at least one provider switch. You can configure the customer-edge interfaces on the PE switches of the MPLS network as either circuit cross-connect (CCC) or using IP over MPLS.

The main differences between configuring IP over MPLS and configuring MPLS over CCC are that for IP over MLPS you configure the customer-edge interfaces to belong to **family inet** rather than **family ccc** and you configure a static route for the label-switched path (LSP). The configuration of the provider switch is the same regardless of whether the PE switches are configured for MPLS over CCC or IP over MPLS. See “Configuring MPLS on Provider Switches (CLI Procedure)” on page 3102.

This topic describes how to configure an ingress PE switch and an egress PE switch for IP over MPLS:

1. Configuring the Ingress PE Switch on page 3108
2. Configuring the Egress PE Switch on page 3109

Configuring the Ingress PE Switch

To configure the ingress PE switch:

1. Configure OSPF (or IS-IS) on the loopback (or switch address) and core interfaces:

```
[edit protocols]
user@switch# set ospf area 0.0.0.0 interface lo0.0
user@switch# set ospf area 0.0.0.0 interface ge-0/0/5.0
user@switch# set ospf area 0.0.0.0 interface ge-0/0/6.0
```

2. Enable traffic engineering for the routing protocol:

```
[edit protocols]
user@switch# set ospf traffic-engineering
```

3. Configure an IP address for the loopback interface and for the core interfaces:

```
[edit]
user@switch# set interfaces lo0 unit 0 family inet address 100.100.100.100/32
user@switch# set interfaces ge-0/0/5 unit 0 family inet address 10.1.5.1/24
user@switch# set interfaces ge-0/0/6 unit 0 family inet address 10.1.6.1/24
```

4. Configure MPLS on the core interfaces:

```
[edit protocols]
user@switch# set mpls interface ge-0/0/5.0
user@switch# set mpls interface ge-0/0/6.0
```

5. Configure **family mpls** on the logical units of the core interfaces, thereby identifying the interfaces that will be used for forwarding MPLS packets:

```
[edit]
user@switch# set interfaces ge-0/0/5 unit 0 family mpls
user@switch# set interfaces ge-0/0/6 unit 0 family mpls
```

6. Configure RSVP on the loopback interface and the core interfaces:

```
[edit protocols]
user@switch# set rsvp interface lo0.0
user@switch# set rsvp interface ge-0/0/5.0
user@switch# set rsvp interface ge-0/0/6.0
```

7. Configure a customer-edge interface as a Layer 3 routed interface, specifying an IP address:

```
[edit]
user@switch# set interfaces ge-2/0/3 unit 0 family inet 121.121.121.1/16
```

8. Configure this Layer 3 customer-edge interface for the routing protocol:

```
[edit]
user@switch# set protocols ospf area 0.0.0.0 interface ge-2/0/3.0
```

9. Configure an LSP on the ingress PE switch (**100.100.100.100**) to send IP packets over MPLS to the egress PE switch (**208.208.208.208**):

```
[edit protocols mpls]
user@switch# set label-switched-path ip_lspjavae_29 from 100.100.100.100
user@switch# set label-switched-path ip_lspjavae_29 to 208.208.208.208
```

10. Disable constrained-path LSP computation for this LSP:

```
[edit protocols mpls]
user@switch# set label-switched-path ip_lspjavae_29 no-cspf
```

11. Configure a static route from the ingress PE switch to the egress PE switch, thereby indicating to the routing protocol that the packets will be forwarded over the MPLS LSP that has been set up to that destination:

```
[edit]
user@switch# set routing-options static route 2.2.2.0/24 next-hop 100.100.100.100
user@switch# set routing-options static route 2.2.2.0/24 resolve
```

Configuring the Egress PE Switch

To configure the egress PE switch:

1. Configure OSPF (or IS-IS) on the loopback interface (or switch address) and core interfaces:

```
[edit protocols]
user@switch# set ospf area 0.0.0.0 interface lo0.0
user@switch# set ospf area 0.0.0.0 interface ge-0/0/5.0
user@switch# set ospf area 0.0.0.0 interface ge-0/0/6.0
```

2. Enable traffic engineering for the routing protocol:

```
[edit protocols]
user@switch# set ospf traffic-engineering
```

3. Configure an IP address for the loopback interface and for the core interfaces:

```
[edit]
user@switch# set interfaces lo0 unit 0 family inet address 208.208.208.208/32
user@switch# set interfaces ge-0/0/5 unit 0 family inet address 10.1.20.1/24
user@switch# set interfaces ge-0/0/6 unit 0 family inet address 10.1.21.1/24
```

4. Configure MPLS on the core interfaces:

```
[edit protocols]
user@switch# set mpls interface ge-0/0/5.0
user@switch# set mpls interface ge-0/0/6.0
```

5. Configure **family mpls** on the logical units of the core interfaces, thereby identifying the interfaces that will be used for forwarding MPLS packets:

```
[edit]
user@switch# set interfaces ge-0/0/5 unit 0 family mpls
user@switch# set interfaces ge-0/0/6 unit 0 family mpls
```

6. Configure RSVP on the loopback interface and the core interfaces:

```
[edit protocols]
user@switch# set rsvp interface lo0.0
user@switch# set rsvp interface ge-0/0/5.0
user@switch# set rsvp interface ge-0/0/6.0
```

7. Configure a customer-edge interface as a Layer 3 routed interface, specifying an IP address:

```
[edit]
user@switch# set interfaces ge-2/0/3 unit 0 family inet address 2.2.2.1/16
```

8. Configure this Layer 3 customer-edge interface for the routing protocol:

```
[edit]
user@switch# set protocols ospf area 0.0.0 interface ge-2/0/3
```

9. Configure an LSP on the egress PE switch (**208.208.208.208**) to send IP packets over MPLS to the ingress PE switch (**100.100.100.100**):

```
[edit protocols mpls]
user@switch# set label-switched-path ip_lsp29_javae from 208.208.208.208
user@switch# set label-switched-path ip_lspjavae_29 to 100.100.100.100
```

10. Disable constrained-path LSP computation for this LSP:

```
[edit protocols mpls]
user@switch# set label-switched-path ip_lsp29_javae no-cspf
```

11. Configure a static route from the ingress PE switch to the egress PE switch, thereby indicating to the routing protocol that the packets will be forwarded over the MPLS LSP that has been set up to that destination:

```
[edit]
user@switch# set routing-options static route 121.121.121.0/24 next-hop 208.208.208.208
user@switch# set routing-options static route 121.121.121.0/24 resolve
```

Related Documentation

- Example: Configuring MPLS on J-EX Series Switches on page 3071
- Configuring MPLS on Provider Switches (CLI Procedure) on page 3102
- Configuring an OSPF Network (J-Web Procedure) on page 1435
- Verifying That MPLS Is Working Correctly on page 3115
- Understanding Junos OS MPLS Components for J-EX Series Switches on page 3059

Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect (CLI Procedure)

Junos OS MPLS for J-EX Series switches supports Layer 2 protocols and Layer 2 virtual private networks (VPNs). You can configure MPLS on J-EX Series switches to increase transport efficiency in your network. MPLS services can be used to connect various sites to a backbone network and to ensure better performance for low-latency applications such as VoIP and other business-critical functions.

To implement MPLS on J-EX Series switches, you must configure two provider edge (PE) switches—an ingress PE switch and an egress PE switch— and at least one provider switch. You can configure the customer-edge interfaces on the PE switches of the MPLS network as either circuit cross-connect (CCC) or using MPLS over IP. For information on configuring MPLS over IP, see "Configuring MPLS on Provider Edge Switches Using MPLS Over IP (CLI Procedure)" on page 3107. You can also use class of service (CoS) on MPLS networks. For further information, see "Understanding Using CoS with MPLS Networks on J-EX Series Switches" on page 2880.

This topic describes configuring PE switches using a circuit cross-connect (CCC). The customer-edge interface can be either a simple interface or a tagged VLAN interface. In both cases, you configure the logical unit of the customer-edge interface to belong to **family ccc** and you must configure an association between that interface and two label-switched paths (LSPs)—one for transmitting MPLS packets to the remote PE and the other for receiving MPLS packets from the remote PE.

The following guidelines apply to CCC configurations:

- When an interface is configured to belong to **family ccc**, it cannot belong to any other family.
- You can send any kind of traffic over a CCC, including nonstandard bridge protocol data units (BPDUs) generated by other vendors' equipment.

If you are configuring a CCC on a tagged VLAN interface, you must explicitly enable VLAN tagging and specify a VLAN ID.

NOTE: The VLAN tag ID cannot be configured on logical interface unit 0. The logical unit number must be 1 or higher.

This procedure shows how to set up two CCCs:

- If you are configuring a CCC on a simple interface (**ge-0/0/1**), you do not need to enable VLAN tagging or specify a VLAN ID.
- If you are configuring a CCC on a tagged VLAN interface (**ge-0/0/2**), include all the steps in this procedure.

To configure a PE switch, complete the following tasks. When you have completed configuring one PE switch, perform the same tasks on the other PE switch:

1. Configure OSPF on the loopback (or switch address) and core interfaces:

```
[edit protocols]
user@switch# set ospf area 0.0.0.0 interface lo0.0
user@switch# set ospf area 0.0.0.0 interface ge-0/0/5.0
user@switch# set ospf area 0.0.0.0 interface ge-0/0/6.0
user@switch# set ospf area 0.0.0.0 interface ae0
```

2. Enable traffic engineering for the routing protocol on both PE switches:

```
[edit protocols]
user@switch# set ospf traffic-engineering
```

3. Configure an IP address for the loopback interface and for the core interfaces:

```
[edit]
user@switch# set interfaces lo0 unit 0 family inet address 127.1.1.1/32
user@switch# set interfaces ge-0/0/5 unit 0 family inet address 10.1.5.1/24
user@switch# set interfaces ge-0/0/6 unit 0 family inet address 10.1.6.1/24
user@switch# set interfaces ae0 unit 0 family inet address 10.1.9.1/24
```

4. Enable MPLS and define the LSP:

```
[edit protocols]
user@switch# set mpls label-switched-path lsp_to_pe2_ge1 to 127.1.1.3
```

TIP: lsp_to_pe2_ge1 is the LSP name. You will need to use the specified name again when configuring the CCC.

5. Configure MPLS on the core interfaces:

```
[edit protocols]
user@switch# set mpls interface ge-0/0/5.0
user@switch# set mpls interface ge-0/0/6.0
user@switch# set mpls interface ae0
```

6. Configure RSVP on the loopback interface and the core interfaces:

```
[edit protocols]
user@switch# set rsvp interface lo0.0
user@switch# set rsvp interface ge-0/0/5.0
user@switch# set rsvp interface ge-0/0/6.0
user@switch# set rsvp interface ae0
```

7. Configure **family mpls** on the logical units of the core interfaces:

```
[edit]
user@switch# set interfaces ge-0/0/5 unit 0 family mpls
user@switch# set interfaces ge-0/0/6 unit 0 family mpls
user@switch# set interfaces ae0 unit 0 family mpls
```

NOTE: You can enable family mpls on either individual interfaces or aggregated Ethernet interfaces. You cannot enable it on tagged VLAN interfaces.

8. If you are configuring a CCC on a tagged VLAN interface, enable VLAN tagging on the customer-edge interface **ge-0/0/2** of the local PE switch:

```
[edit interfaces ge-0/0/2]
user@switch# set vlan-tagging
```

   If you are configuring a CCC on a simple interface (**ge-0/0/1**), omit this step.

9. If you are configuring a CCC on a tagged VLAN interface, configure the logical unit of the customer-edge interface with a VLAN ID:

```
[edit interfaces ge-0/0/2 unit 1]
user@switch# set vlan-id 100
```

   If you are configuring a CCC on a simple interface (**ge-0/0/1**), omit this step.

10. Configure the logical unit of the customer-edge interface to belong to **family ccc**:

```
[edit interfaces ge-0/0/1 unit 0]
user@switch# set family ccc

[edit interfaces ge-0/0/2 unit 1]
user@switch# set family ccc
```

11. Associate the CCC interface with two LSPs, one for transmitting MPLS packets and the other for receiving MPLS packets:

```
[edit protocols]
user@switch# set connections remote-interface-switch ge-1–to-pe2 interface ge-0/0/1.0
user@switch# set connections remote-interface-switch ge-1–to-pe2 transmit-lsp
lsp_to_pe2_ge1
user@switch# set connections remote-interface-switch ge-1–to-pe2 receive-lsp
lsp_to_pe1_ge1

[edit protocols]
user@switch# set connections remote-interface-switch ge-1–to-pe2 interface ge-0/0/2.1
user@switch# set connections remote-interface-switch ge-1–to-pe2 transmit-lsp
lsp_to_pe2_ge1
user@switch# set connections remote-interface-switch ge-1–to-pe2 receive-lsp
lsp_to_pe1_ge1
```

TIP: The transmit-lsp option specifies the LSP name that was configured on PE-1 (the local PE switch) by the label-switched-path statement within the protocols mpls stanza.
The receive-lsp option specifies the LSP name that was configured on PE-2 (the remote PE switch) by the label-switched-path statement within the protocols mpls stanza.

When you have completed configuring one PE switch, follow the same procedures to configure the other PE switch.

Related Documentation

- Example: Configuring MPLS on J-EX Series Switches on page 3071
- Configuring MPLS on Provider Switches (CLI Procedure) on page 3102

- Configuring MPLS on Provider Edge Switches Using MPLS Over IP (CLI Procedure) on page 3107
- Configuring an OSPF Network (J-Web Procedure) on page 1435
- Verifying That MPLS Is Working Correctly on page 3115
- Understanding Junos OS MPLS Components for J-EX Series Switches on page 3059

CHAPTER 123

Verifying MPLS

Verifying That MPLS Is Working Correctly

To verify that MPLS is working correctly on J-EX Series switches, perform the following tasks:

Verifying the Physical Layer on the Switches

Purpose Verify that the interfaces are up. Perform this verification task on each of the switches.

Action

```
user@switch> show interfaces ge- terse

Interface               Admin Link Proto    Local                 Remote
ge-0/0/0                up    up
ge-0/0/0.0              up    up
ge-0/0/1.0              up    up   ccc
ge-0/0/2.0              up    up   ccc
ge-0/0/3.0              up    up   eth-switch
ge-0/0/4.0              up    up   eth-switch
ge-0/0/5.0              up    up   inet      10.1.5.1/24
            mpls
ge-0/0/6.0              up    up   inet      10.1.6.1/24
            mpls
```

Meaning The **show interfaces terse** command displays status information about the Gigabit Ethernet interfaces on the switch. This output verifies that the interfaces are **up**. The output for the protocol family (**Proto** column) shows that interfaces **ge-0/0/1.0** and **ge-0/0/2.0** are configured as circuit cross-connect. The Local and Remote columns do not display

IP addresses, because the **inet family** is not configured for CCC interfaces. The output for the protocol family of the core interfaces (**ge-0/0/0.5** and **ge-0/0/0.6**), shows that these interfaces are configured as both **inet** and **mpls**. The **Local** column for the core interfaces shows the IP address configured for these interfaces.

Verifying the Routing Protocol

Purpose Verify the state of the configured routing protocol. You should perform this verification task on each of the switches. The state should be **Full**. If you have configured OSPF as the routing protocol, use the **show ospf neighbor** command to verify that the routing protocol is communicating with the switch neighbors. If you have configured IS-IS as the routing protocol, use the **show isis adjacency** command to verify that the routing protocol is communicating with the switch neighbors.

Action

```
user@switch> show ospf neighbor

Address          Interface              State     ID               Pri  Dead
127.1.1.2          ge–0/0/5               Full      10.10.10.10      128    39
```

Meaning The **show ospf neighbor** command displays the status of the routing protocol that has been configured on this switch. The output shows that the state is **full**, meaning that the routing protocol is operating correctly—that is, hello packets are being exchanged between directly connected neighbors. For additional information on checking and monitoring routing protocols, see the *Junos OS Routing Protocols and Policies Command Reference* at http://www.juniper.net/techpubs/software/junos/ .

Verifying the Core Interfaces Being Used for the MPLS Traffic

Purpose Verify that the state of the MPLS interface is **Up**. You should perform this verification task on each of the switches.

Action

```
user@switch> show mpls interface

Interface        State       Administrative groups
ge–0/05            Up          <none>
```

Meaning The **show mpls interface** command displays the status of the core interfaces that have been configured to belong to **family mpls**. This output shows that the interface configured to belong to **family mpls** is **up**.

Verifying RSVP

Purpose Verify the state of the RSVP session. You should perform this verification task on each of the switches.

Action

```
user@switch> show rsvp session

Ingress RSVP: 1 sessions
To              From            State   Rt Style Labelin Labelout LSPname
127.1.1.3         127.1.1.1         Up       0  1 FF       -   300064
lsp_to_pe2_ge1
Total 1 displayed, Up 1, Down 0
```

```
Egress RSVP: 1 sessions
To              From            State   Rt Style Labelin Labelout LSPname
127.1.1.1         127.1.1.3         Up       0  1 FF  299968        -
lsp_to_pe1_ge1
Total 1 displayed, Up 1, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

Meaning This output confirms that the RSVP sessions are **Up**.

Verifying the Assignment of Interfaces for MPLS Label Operations

Purpose Verify which interface is being used as the beginning of the CCC and which interface is being used to push the MPLS packet to the next hop. You should perform this task only on the provider edge switches.

Action

```
user@switch> show route forwarding-table family mpls
MPLS:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    dscd    50     1
0                  user     0                    recv    49     3
1                  user     0                    recv    49     3
2                  user     0                    recv    49     3
299776             user     0                    Pop    541     2 ge-0/0/1.0
ge-0/0/1.0  (CCC)  user     0 127.1.2.1         Push 299792   540 2   ge-0/0/5.0
```

Meaning This output shows that CCC has been set up on interface **ge-0/0/1.0**. The switch receives ingress traffic on **ge-0/0/1.0** with label **299776**. It pops that label and swaps it to label **299792**, which it pushes out on interface **ge-0/0/5.0**.

Verifying the Status of the CCC

Purpose Verify the status of the CCC. You should perform this task only on the provider edge switches.

Action

```
user@switch> show connections
CCC and TCC connections [Link Monitoring On]
Legend for status (St)              Legend for connection types
UN -- uninitialized                 if-sw:  interface switching
NP -- not present                   rmt-if: remote interface switching
WE -- wrong encapsulation           lsp-sw: LSP switching
DS -- disabled                      tx-p2mp-sw: transmit P2MP switching
Dn -- down                          rx-p2mp-sw: receive P2MP switching
-> -- only outbound conn is up
 <- -- only inbound  conn is up      Legend for circuit types
Up -- operational                   intf -- interface
RmtDn -- remote CCC down            tlsp -- transmit LSP
Restart -- restarting               rlsp -- receive LSP


Connection/Circuit               Type        St      Time last up     # Up trans
ge1-to-pe2                        rmt-if      Up      Feb 17 05:00:09  1
  ge-0/0/1.0                       intf       Up
```

```
  lsp_to_pe1_ge1                    tlsp     Up
  lsp_to_pe2_ge1                    rlsp     Up
```

Meaning The **show connections** command displays the status of the CCC connections. This output verifies that the CCC interface and its associated transmit and receive LSPs are **Up**.

Related Documentation

- Configuring MPLS on Provider Edge Switches (CLI Procedure)
- Configuring MPLS on Provider Switches (CLI Procedure) on page 3102

Verifying Path Protection in an MPLS Network

To verify that path protection is working correctly on J-EX Series switches, perform the following tasks:

1. Verifying the Primary Path on page 3118
2. Verifying the RSVP-Enabled Interfaces on page 3119
3. Verifying a Secondary Path on page 3119

Verifying the Primary Path

Purpose Verify that the primary path is operational.

Action

```
user@switch> show mpls lsp extensive ingress

Ingress LSP: 2 sessions

127.1.8.8
  From: 127.1.9.9, State: Up, ActiveRoute: 0, LSPname: lsp_to_240
  ActivePath: primary_path_lsp_to_240 (primary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   primary_path_lsp_to_240 State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
          Exclude: red
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 2)
 10.3.3.2 S 10.3.4.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
 20=Node-ID):
          10.3.3.2 10.3.4.2
    6 Mar 11 23:58:01.684 Selected as active path: due to 'primary'
    5 Mar 11 23:57:00.750 Record Route:  10.3.3.2 10.3.4.2
    4 Mar 11 23:57:00.750 Up
    3 Mar 11 23:57:00.595 Originate Call
    2 Mar 11 23:57:00.595 CSPF: computation result accepted  10.3.3.2 10.3.4.2
    1 Mar 11 23:56:31.135 CSPF failed: no route toward 10.3.2.2[25 times]
Standby   secondary_path_lsp_to_240 State: Up
Standby   secondary_path_lsp_to_240 State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 1)
 10.3.5.2 S
    Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt
 20=Node-ID):
```

```
        10.3.5.2
     7 Mar 11 23:58:01.684 Deselected as active: due to 'primary'
     6 Mar 11 23:46:17.298 Selected as active path
     5 Mar 11 23:46:17.295 Record Route:  5.5.5.2
     4 Mar 11 23:46:17.287 Up
     3 Mar 11 23:46:16.760 Originate Call
     2 Mar 11 23:46:16.760 CSPF: computation result accepted  10.3.5.2
     1 Mar 11 23:45:48.095 CSPF failed: no route toward 10.5.5.5[2 times]
  Created: Wed Mar 11 23:44:37 2009
 [Output truncated]
```

Meaning As indicated by the **ActivePath** in the output, the LSP **primary_path_lsp_to_240** is active.

Verifying the RSVP-Enabled Interfaces

Purpose Verify the status of Resource Reservation Protocol (RSVP)-enabled interfaces and packet statistics.

Action

```
user@switch> show rsvp interfaces

RSVP interface: 1 active
                    Active Subscr- Static      Available   Reserved    Highwater
Interface   State resv   iption  BW          BW          BW          mark
ge-0/0/20.0 Up         2   100%  1000Mbps    1000Mbps    0bps        0bps
```

Meaning This output verifies that RSVP is enabled and operational on interface **ge-0/0/20.0**.

Verifying a Secondary Path

Purpose Verify that a secondary path is established.

Action Deactivate a switch that is critical to the primary path and then issue the following command:

```
user@switch> show mpls lsp extensive

Ingress LSP: 1 sessions

127.0.0.8
  From: 127.0.0.1, State: Up, ActiveRoute: 0, LSPname: lsp_to_240
  ActivePath: secondary_path_lsp_to_240 (secondary)
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
  Primary   primary_path_lsp_to_240 State: Dn
    Priorities: 7 0
    SmartOptimizeTimer: 180
          Exclude: red
    Will be enqueued for recomputation in 8 second(s).
   51 Mar  8 12:23:31.268 CSPF failed: no route toward 127.0.0.11[11420 times]
   50 Mar  4 15:35:25.610 Clear Call: CSPF computation failed
   49 Mar  4 15:35:25.610 CSPF: link down/deleted:
127.0.0.2(127.0.0.1:0)(127.0.0.1)->
0.0.0.0(127.0.0.20:0)(127.0.0.20)
   48 Mar  4 15:35:25.576 Deselected as active
   47 Mar  4 15:35:25.550 No Route toward dest
   46 Mar  4 15:35:25.550 ?????
   45 Mar  4 15:35:25.549 127.0.0.12: Down
```

```
   44 Mar  4 15:33:29.839 Selected as active path
   43 Mar  4 15:33:29.837 Record Route:  127.0.0.20 127.0.0.40
   42 Mar  4 15:33:29.835 Up
   41 Mar  4 15:33:29.756 Originate Call
  40 Mar  4 15:33:29.756 CSPF: computation result accepted  127.0.0.20 127.0.0.40

   39 Mar  4 15:33:00.395 CSPF failed: no route toward 127.0.0.11[7 times]
   38 Mar  4 15:30:31.412 Clear Call: CSPF computation failed
   37 Mar  4 15:30:31.412 CSPF: link down/deleted:
127.0.0.2(127.0.0.1:0)(127.0.0.1)->
0.0.0.0(127.0.0.20:0)(127.0.0.20)
   36 Mar  4 15:30:31.379 Deselected as active
   35 Mar  4 15:30:31.350 No Route toward dest
   34 Mar  4 15:30:31.350 ?????
   33 Mar  4 15:30:31.349 127.0.0.12: Down
   32 Mar  4 15:29:05.802 Selected as active path
   31 Mar  4 15:29:05.801 Record Route:  127.0.0.20 127.0.0.40
   30 Mar  4 15:29:05.801 Up
   29 Mar  4 15:29:05.686 Originate Call
  28 Mar  4 15:29:05.686 CSPF: computation result accepted  127.0.0.20 127.0.0.40

   27 Mar  4 15:28:35.852 CSPF failed: no route toward 127.0.0.11[132 times]
   26 Mar  4 14:25:12.113 Clear Call: CSPF computation failed
   25 Mar  4 14:25:12.113 CSPF: link down/deleted:
0.0.0.0(127.0.0.20:0)(127.0.0.20)->
0.0.0.0(10.10.10.10:0)(10.10.10.10)
 *Standby   secondary_path_lsp_to_240 State: Up
    Priorities: 7 0
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 1)
 [Output truncated]
```

Meaning As indicated by the **ActivePath** in the output, the LSP **secondary_path_lsp_to_240** is active.

Related Documentation

- Configuring Path Protection in an MPLS Network (CLI Procedure) on page 3097
- Understanding MPLS and Path Protection on J-EX Series Switches on page 3063
- For information on the **show mpls lsp** and **show rvsp interfaces** commands, see the *Junos OS MPLS Applications Configuration Guide* at http://www.juniper.net/techpubs/software/junos/junos95/index.html.

CHAPTER 124

Configuration Statements for MPLS

[edit protocols] Configuration Statement Hierarchy

```
protocols {
  connections {
    remote-interface-switch connection-name {
      interface interface-name.unit-number;
      transmit-lsp label-switched-path;
      receive-lsp label-switched-path;
    }
  }
  dot1x {
    authenticator {
      authentication-profile-name profile-name;
      interface (all | [ interface-names ]) {
        disable;
        guest-vlan ( vlan-id | vlan-name);
        mac-radius <restrict>;
        maximum-requests number;
        no-reauthentication;
        quiet-period seconds;
        reauthentication {
          interval seconds;
        }
        retries number;
        server-fail (deny | permit | use-cache | vlan-id | vlan-name);
        server-reject-vlan (vlan-id | vlan-name);
        server-timeout seconds;
        supplicant (multiple | single | single-secure);
        supplicant-timeout seconds;
        transmit-period seconds;
      }
      static mac-address {
        interface interface-name;
        vlan-assignment (vlan-id |vlan-name);
      }
  }
  gvrp {
    <enable | disable>;
    interface (all | [interface-name]) {
```

```
      disable;
    }
    join-timer millseconds;
    leave-timer milliseconds;
    leaveall-timer milliseconds;
  }
  igmp-snooping {
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>
        <match regex>;
      flag flag (detail | disable | receive | send);
    }
    vlan (vlan-id | vlan-number) {
      data-forwarding {
        source {
          groups group-prefix;
        }
        receiver {
          source-vlans vlan-list;
          install ;
        }
      }
      disable {
        interface interface-name
      }
      immediate-leave;
      interface interface-name {
        group-limit limit;
        multicast-router-interface;
        static {
          group ip-address;
        }
      }
      proxy ;
      query-interval seconds;
      query-last-member-interval seconds;
      query-response-interval seconds;
      robust-count number;
    }
  }
  lldp {
    disable;
    advertisement-interval seconds;
    hold-multiplier number;
    interface (all | interface-name) {
      disable;
    }
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>
        <match regex>;
      flag flag (detail | disable | receive | send);
    }
  }
  lldp-med {
    disable;
    fast-start number;
```

```
    interface (all | interface-name) {
      disable;
      location {
        elin number;
        civic-based {
          what number;
          country-code code;
          ca-type {
            number {
              ca-value value;
            }
          }
        }
      }
    }
  }
  mpls {
    interface ( all | interface-name );
    label-switched-path lsp-name to remote-provider-edge-switch;
    path destination {
      <address | hostname> <strict | loose>
    }
  mstp {
    disable;
    bpdu-block-on-edge;
    bridge-priority priority;
    configuration-name name;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
      disable;
      bpdu-timeout-action {
        block;
        alarm;
      }
      cost  cost;
      edge;
      mode mode;
      no-root-port;
      priority priority;
    }
    max-age seconds;
    max-hops hops;
    msti msti-id {
      vlan (vlan-id | vlan-name);
      interface interface-name {
        disable;
        cost  cost;
        edge;
        mode mode;
        priority priority;
      }
    }
    revision-level revision-level;
    traceoptions {
```

```
    file filename <files number > <size size> <no-stamp | world-readable |
      no-world-readable>;
    flag flag;
  }
}
mvrp {
  disable
  interface (all | interface-name) {
    disable;
    join-timer milliseconds;
    leave-timer milliseconds;
    leaveall-timer milliseconds;
    registration (forbidden | normal);
  }
  no-dynamic-vlan;
  traceoptions {
    file filename <files number > <size size> <no-stamp | world-readable |
      no-world-readable>;
    flag flag;
  }
}
oam {
  ethernet{
    connectivity-fault-management {
      action-profile profile-name {
        default-actions {
          interface-down;
        }
      }
      linktrace {
        age (30m | 10m | 1m | 30s | 10s);
        path-database-size path-database-size;
      }
      maintenance-domain domain-name {
        level number;
        mip-half-function (none | default |explicit);
        name-format (character-string | none | dns | mac+2oct);
        maintenance-association ma-name {
          continuity-check {
            hold-interval minutes;
            interval (10m | 10s | 1m | 1s| 100ms);
            loss-threshold number;
          }
          mep mep-id {
            auto-discovery;
            direction down;
            interface interface-name;
            remote-mep mep-id {
              action-profile profile-name;
            }
          }
        }
      }
    }
    link-fault-management {
      action-profile profile-name;
```

```
            action {
              syslog;
              link-down;
            }
            event {
              link-adjacency-loss;
              link-event-rate;
              frame-error count;
              frame-period count;
              frame-period-summary count;
              symbol-period count;
            }
            interface interface-name {
              link-discovery (active | passive);
              pdu-interval interval;
              event-thresholds threshold-value;
              remote-loopback;
              event-thresholds {
                frame-errorcount;
                frame-period count;
                frame-period-summary count;
                symbol-period count;
              }
            }
            negotiation-options {
              allow-remote-loopback;
              no-allow-link-events;
            }
        }
    }
}
rstp {
    disable;
    bpdu-block-on-edge;
    bridge-priority priority;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
        disable;
        bpdu-timeout-action {
            block;
            alarm;
        }
        cost  cost;
        edge;
        mode mode;
        no-root-port;
        priority priority;
    }
    max-age seconds;
    }
    traceoptions {
        file filename <files number > <size size> <no-stamp | world-readable |
            no-world-readable>;
        flag flag;
    }
```

```
}
sflow {
  agent-id
  collector {
    ip-address;
    udp-port port-number;
  }
  disable;
  interfaces interface-name {
    disable;
    polling-interval seconds;
    sample-rate number;
  }
  polling-interval seconds;
  sample-rate number;
  source-ip
}
stp {
  disable;
  bridge-priority priority;
  forward-delay seconds;
  hello-time seconds;
  interface (all | interface-name) {
    disable;
    bpdu-timeout-action {
      block;
      alarm;
    }
    cost  cost;
    edge;
    mode mode;
    no-root-port;
    priority priority;
  }
  max-age seconds;
}
  traceoptions {
    file filename <files number > <size size> <no-stamp | world-readable |
      no-world-readable>;
    flag flag;
  }
vstp {
  bpdu-block-on-edge;
  disable;
  force-version stp;
  vlan (all | vlan-id | vlan-name) {
    bridge-priority priority;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
      bpdu-timeout-action {
        alarm;
        block;
      }
      cost  cost;
      disable;
```

```
            edge;
            mode mode;
            no-root-port;
            priority priority;
          }
          max-age seconds;
          traceoptions {
            file filename <files number > <size size> <no-stamp | world-readable |
              no-world-readable>;
            flag flag;
          }
        }
      }
    }
```

Related Documentation

- 802.1X for J-EX Series Switches Overview on page 2253
- Example: Configure Automatic VLAN Administration Using GVRP on page 1087
- Understanding MAC RADIUS Authentication on J-EX Series Switches
- Understanding Server Fail Fallback and 802.1X Authentication on J-EX Series Switches on page 2258
- IGMP Snooping on J-EX Series Switches Overview on page 2047
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261
- Understanding MSTP for J-EX Series Switches on page 1277
- Understanding Multiple VLAN Registration Protocol (MVRP) on J-EX Series Switches on page 1054
- Understanding Ethernet OAM Connectivity Fault Management for a J-EX Series Switch on page 3463
- Understanding Ethernet OAM Link Fault Management for a J-EX Series Switch on page 3427
- Understanding RSTP for J-EX Series Switches on page 1276
- Understanding STP for J-EX Series Switches on page 1275
- Understanding How to Use sFlow Technology for Network Monitoring on a J-EX Series Switch on page 3283
- Understanding VSTP for J-EX Series Switches on page 1281

connections

Syntax

```
connections {
  remote-interface-switch connection-name {
    interface interface-name.unit-number;
    transmit-lsp label-switched-path;
    receive-lsp label-switched-path;
  }
```

Hierarchy Level [edit protocols]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define the connection between two circuits in a CCC connection.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring MPLS on J-EX Series Switches on page 3071
- Configuring MPLS on Provider Edge Switches (CLI Procedure)
- *Junos OS MPLS Applications Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

exp

Syntax

```
exp classifier-name {
  import (classifier-name | default);
  forwarding-class class-name {
    loss-priority level {
      code-points [ aliases ] [ 3–bit-patterns ];
    }
  }
}
```

Hierarchy Level [edit class-of-service classifiers]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define the experimental bits (EXP) code point mapping that is applied to the MPLS packets.

J-EX Series switches support only one EXP code mapping on the switch (either default or custom). It is applied globally and implicitly to all the MPLS-enabled interfaces on the switch. You cannot bind it to an individual interface and you cannot disable it.

Options *classifier-name*—Name of the classifier.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Understanding Using CoS with MPLS Networks on J-EX Series Switches on page 2880
- Configuring MPLS on Provider Edge Switches Using Circuit Cross-Connect (CLI Procedure) on page 3111
- Configuring MPLS on Provider Edge Switches Using IP Over MPLS (CLI Procedure) on page 3107
- Configuring CoS on Provider Switches of an MPLS Network (CLI Procedure) on page 3106

interface

Syntax interface (all | *interface-name*);

Hierarchy Level [edit protocols mpls]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable MPLS on all interfaces on the switch or on the specified interface.

Default MPLS is disabled.

Options **all**—All interfaces on the switch.

interface-name—Name of an interface:

- Aggregated Ethernet—**ae*x***
- Gigabit Ethernet—**ge-*fpc/pic/port***

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring MPLS on J-EX Series Switches on page 3071
- Configuring MPLS on Provider Edge Switches (CLI Procedure)
- Configuring MPLS on Provider Switches (CLI Procedure) on page 3102

label-switched-path

Syntax label-switched-path *lsp-name* to *remote-provider-edge-switch*;

Hierarchy Level [edit protocols mpls]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define a label-switched path (LSP) to the remote provider edge switch to use for MPLS traffic. You must specify this statement on the provider edge switch.

Options

lsp-name —Name that identifies the LSP. The name can be up to 32 characters and can contain letters, digits, periods, and hyphens. To include other characters, enclose the name in quotation marks. The name must be unique on the ingress switch.

remote-provider-edge-switch —Either the loopback address or the switch address.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring MPLS on J-EX Series Switches on page 3071
- Configuring MPLS on Provider Edge Switches (CLI Procedure)
- *Junos OS MPLS Applications Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

mpls

Syntax

```
mpls {
  interface ( all | interface-name );
  label-switched-path lsp-name to remote-provider-edge-switch;
  path destination {
    <address | hostname> <strict | loose>
  }
}
```

Hierarchy Level [edit protocols]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable MPLS on the switch.

The remaining statements are explained separately.

Default MPLS is disabled.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring MPLS on J-EX Series Switches on page 3071
- Configuring MPLS on Provider Edge Switches (CLI Procedure)
- Configuring MPLS on Provider Switches (CLI Procedure) on page 3102
- *Junos OS MPLS Applications Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

path

Syntax

```
path destination {
  <address | hostname> <strict | loose>
}
```

Hierarchy Level [edit protocols mpls]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure path protection on your MPLS network.

Options

destination —Name of a label switched path (LSP). In addition to specifying the name of the configured LSP, you can include some other designation such as **primary-path**.

address —(Optional) IP address of each transit switch (or the IP address of the loopback interface on the switch) in the LSP. If you want to control exactly which switches are selected for the LSP, specify the address or hostname of each transit switch. Specify the addresses in order, starting with the first provider (transit) switch, and continuing sequentially along the path until reaching the egress provider edge switch.

Default: If you do not specify the addresses or hostnames of any switches, the LSP is calculated by the switch.

hostname —(Optional) See ***address*** .

Default: If you do not specify the addresses or hostnames of any switches, the LSP is calculated by the switch.

loose—(Optional) Indicates that the next address in the **path** statement is a loose link. This means that the LSP can traverse through other switches before reaching this switch.

Default: **strict**

strict—(Optional) Indicates that the LSP must go to the next address specified in the **path** statement without traversing other switches. This is the default.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Path Protection in an MPLS Network (CLI Procedure) on page 3097

policing

Syntax policing (filter *filter-name* | no-automatic-policing);

Hierarchy Level [edit protocols mpls label-switched-path *lsp-name*]
[edit interfaces *interface-id* unit *number-of-logical-unit* family inet address *ip-address*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Apply a rate-limiting policer as the specified policing filter:

- To the LSP for MPLS over CCC.
- To the customer-edge interface for IP over MPLS.

Options **filter *filter-name***—Specify the name of the policing filter.

no-automatic-policing—Disable automatic policing on this LSP.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- **policer on page 2826**
- Configuring Policers to Control Traffic Rates (CLI Procedure) on page 2788
- Configuring CoS on MPLS Provider Edge Switch Using Circuit Cross-Connect (CLI Procedure) on page 2932
- Configuring CoS on MPLS Provider Edge Switch Using IP Over MPLS (CLI Procedure) on page 2931

primary

Syntax primary *path-name*;

Hierarchy Level [edit protocols **mpls label-switched-path** *lsp-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the primary path to use for a label switched path (LSP). You can configure only one primary path.

Options ***path-name*** —Name of the primary path that you created with the **path** statement.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Path Protection in an MPLS Network (CLI Procedure) on page 3097

remote-interface-switch

Syntax

```
remote-interface-switch connection-name {
    interface interface-name.unit-number;
    receive-lsp label-switched-path;
    transmit-lsp label-switched-path;
}
```

Hierarchy Level [edit protocols **connections**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure MPLS LSP tunnel cross-connects. This makes an association between a CCC interface and two LSPs, one for transmitting MPLS packets from the local provider edge switch to the remote provider edge switch and the other for receiving MPLS packets on the local provider edge switch from the remote provider edge switch.

Options

connection-name —Connection name.

interface *interface-name.unit-number* —Interface name. Include the logical portion of the name, which corresponds to the logical unit number of the CCC interface.

receive-lsp *label-switched-path* —Name of the LSP from the connection's source. This LSP name was specified by the **label-switched-path** statement on the remote provider edge switch in the **protocols mpls** stanza.

transmit-lsp *label-switched-path* —Name of the LSP to the connection's destination. This LSP name was specified by the **label-switched-path** statement on the local provider edge switch in the **protocols mpls** stanza.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring MPLS on J-EX Series Switches on page 3071
- Configuring MPLS on Provider Edge Switches (CLI Procedure)
- *Junos OS MPLS Applications Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

revert-timer

Syntax revert-timer *seconds*;

Hierarchy Level [edit protocols **mpls**],
[edit protocols mpls **label-switched-path** *lsp-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the amount of time that a label switched path (LSP) must wait before traffic reverts to a primary path. If during this time the primary path experiences any connectivity problem or stability problem, the timer is restarted.

If you have configured a value of 0 seconds for the **revert-timer** statement and traffic is switched to the secondary path, the traffic remains on that path indefinitely. It is never switched back to the primary path unless you intervene.

Default 60 seconds

Options ***seconds*** —Value in seconds.

Range: 0 through 65,535 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Path Protection in an MPLS Network (CLI Procedure) on page 3097

rsvp

Syntax rsvp;

Hierarchy Level [edit protocols]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable Resource Reservation Protocol (RSVP) signaling. The primary purpose of RSVP in the Junos OS for J-EX Series switches is to support dynamic signaling within label switched paths (LSPs).

Default RSVP is disabled.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring MPLS on J-EX Series Switches on page 3071
- Configuring MPLS on Provider Edge Switches (CLI Procedure)
- Configuring MPLS on Provider Switches (CLI Procedure) on page 3102
- *Junos OS MPLS Applications Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

secondary

Syntax

```
secondary path-name {
    standby;
{
```

Hierarchy Level [edit protocols **mpls label-switched-path** *lsp-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify one or more secondary paths to use for the label switched path (LSP). You can configure more than one secondary path. All secondary paths are equal, and the first one that is available is chosen.

Options ***path-name*** —Name of a secondary path that you created with the **path** statement.

The remaining statement is explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Configuring Path Protection in an MPLS Network (CLI Procedure) on page 3097

standby

Syntax standby;

Hierarchy Level [edit protocols **mpls label-switched-path** *lsp-name* **secondary** *path-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable the path to remain up at all times to provide instant switchover if connectivity problems occur.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Configuring Path Protection in an MPLS Network (CLI Procedure) on page 3097

traffic-engineering

Syntax traffic-engineering;

Hierarchy Level [edit protocols ospf | isis]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable the traffic engineering features of the specified routing protocol.

Default Traffic engineering is disabled.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring MPLS on J-EX Series Switches on page 3071
- Configuring MPLS on Provider Edge Switches (CLI Procedure)
- Configuring MPLS on Provider Switches (CLI Procedure) on page 3102
- Configuring an OSPF Network (J-Web Procedure) on page 1435
- *Junos OS MPLS Applications Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

CHAPTER 125

Operational Mode Commands for MPLS

clear mpls lsp

Syntax

clear mpls lsp
<autobandwidth>
<logical-system (all | *logical-system-name*)>
<name *name*>
<optimize | optimize-aggressive>
<path *regular-expression*>
<statistics>

Syntax (J-EX Series Switch)

clear mpls lsp
<autobandwidth>
<name *name*>
<optimize | optimize-aggressive>
<path *regular-expression*>
<statistics>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Release the routes and states associated with MPLS label-switched paths (LSPs), and start new LSPs.

CAUTION: This command disconnects existing Resource Reservation Protocol (RSVP) sessions on the ingress routing device. If there is a time lag between the old path being torn down and the new path being set up, this command might impact traffic traveling along the LSPs.

Options

none—Reset and restart all LSPs that originated from this routing device; that is, all LSPs for which this routing device is the ingress routing device. Depending on the number of LSPs involved, it might take a while to restart all the LSPs.

autobandwidth—(Optional) Clear LSP autobandwidth counters.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

name *name*—(Optional) Reset and restart the specified LSP or group of LSPs. You can include wildcard characters in the interface name, as described in the *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.

optimize | optimize-aggressive—(Optional) Run nonpreemptive optimization or aggressive optimization computation now.

path *regular-expression*—(Optional) Clear the specific LSP path matching the specified regular expression.

statistics—(Optional) Clear LSP statistics.

Required Privilege Level clear

Related Documentation
- show mpls lsp on page 3189
- show rsvp session on page 3221

List of Sample Output clear mpls lsp on page 3141

Output Fields When you enter this command, you are provided feedback on the status of your request.

clear mpls lsp

```
user@host> clear mpls lsp
```

clear rsvp session

Syntax

clear rsvp session
<connection-source *address*>
<connection-destination *address*>
<gracefully>
<logical-system (all | *logical-system-name*)>
<lsp-id *identifier*>
<name *name*>
<optimize-fast-reroute>
<tunnel-id *identifier*>

Syntax (J-EX Series Switch)

clear rsvp session
<connection-source *address*>
<connection-destination *address*>
<gracefully>
<lsp-id *identifier*>
<name *name*>
<optimize-fast-reroute>
<tunnel-id *identifier*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Reset and restart Resource Reservation Protocol (RSVP) sessions.

Options

none—Reset and restart all RSVP sessions for which this routing device is the ingress, transit, or egress routing device.

connection-source *address*—(Optional) Source address for GMPLS and MPLS LSPs from the RSVP sender template.

connection-destination *address*—(Optional) Destination address for GMPLS and MPLS LSPs from the RSVP sender template.

gracefully—(Optional) Gracefully reset an RSVP session for a nonpacket LSP in two passes. In the first pass, the Admin-Status object is signaled along the path to the other endpoint of the RSVP session. In the second pass, the path used by the RSVP session is torn down. This option can only be used on the ingress or egress routing device of the RSVP session and is only valid for nonpacket LSPs.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

lsp-id *identifier*—(Optional) LSP identifier (source port) for the RSVP sender template.

name *name*—(Optional) Reset and restart the specified RSVP session.

optimize-fast-reroute—(Optional) Begin fast reroute optimization.

tunnel-id *identifier*—(Optional) Tunnel identifier (destination port) for the RSVP session.

Required Privilege Level clear

Related Documentation
- clear mpls lsp on page 3140
- show rsvp session on page 3221

List of Sample Output clear rsvp session on page 3143

Output Fields When you enter this command, you are provided feedback on the status of your request.

clear rsvp session

```
user@host> clear rsvp session
```

clear rsvp statistics

Syntax
clear rsvp statistics
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
clear rsvp statistics

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Clear Resource Reservation Protocol (RSVP) packet and error statistics.

Options
none—Clear RSVP packet and error statistics.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
clear

Related Documentation
- **show rsvp statistics on page 3229**

List of Sample Output
clear rsvp statistics on page 3144

Output Fields
When you enter this command, you are provided feedback on the status of your request.

clear rsvp statistics
```
user@host> clear rsvp statistics
```

ping mpls l2circuit

Syntax ping mpls l2circuit (interface *interface-name* | virtual-circuit *virtual-circuit-id* neighbor *address*)
<count *count*>
<destination *address*>
<detail>
<exp *forwarding-class*>
<logical-system (all | *logical-system-name*)>
<size *bytes*>
<source *source-address*>
<sweep>
<v1>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Check the operability of the MPLS Layer 2 circuit connections. Type Ctrl+c to interrupt a **ping mpls l2circuit** command.

Options

count *count*—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is **1** through **1,000,000**. The default value is **5**.

destination *address*—(Optional) Specify an address other than the default (**127.0.0.1/32**) for the ping echo requests. The address can be anything within the **127/8** subnet.

detail—(Optional) Display detailed information about the echo requests sent and received.

exp *forwarding-class*—(Optional) Value of the forwarding class for the MPLS ping packets.

interface *interface-name*—Ping an interface configured for the Layer 2 circuit on the egress provider edge (PE) router.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on the specified logical system.

size *bytes*—(Optional) Size of the label-switched path (LSP) ping request packet (**96** through **65468** bytes). Packets are 4-byte aligned. For example, If you enter a size of 97, 98, 99, or 100, the router or switch uses a size value of 100 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 96-byte minimum.

source *source-address*—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (**lo.0**).

sweep—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

v1—(Optional) Use the type 9 Layer 2 circuit type, length, and value (TLV).

virtual-circuit *virtual-circuit-id* neighbor *address*—Ping the virtual circuit identifier on the egress PE router or switch and the specified neighbor, testing the integrity of the Layer 2 circuit between the ingress and egress PE routers or switches.

Additional Information

You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the egress PE router or switch (the router or switch receiving the MPLS echo packets) to ping a Layer 2 circuit.

In asymmetric MTU scenarios, the echo response may be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

Required Privilege Level

network

List of Sample Output

ping mpls l2circuit interface on page 3146
ping mpls l2circuit virtual-circuit detail on page 3146

Output Fields

When you enter this command, you are provided feedback on the status of your request. An exclamation point (**!**) indicates that an echo reply was received. A period (**.**) indicates that an echo reply was not received within the timeout period. An **x** indicates that an echo reply was received with an error code. Packets with an error code are not counted in the received packets count. They are accounted for separately.

ping mpls l2circuit interface

```
user@host> ping mpls l2circuit interface so-1/0/0.1
Request for seq 1, to interface 69, labels <100000, 100208>, packet size 100
Reply for seq 1, return code: Egress-ok, time: 0.439 ms
```

ping mpls l2circuit virtual-circuit detail

```
user@host> ping mpls l2circuit virtual-circuit 200 neighbor 10.255.245.122/32 detail
Request for seq 1, to interface 68, labels <100048, 100128>, packet size 100

Reply for seq 1, return code: Egress-ok time: 0.539 ms
```

ping mpls l2vpn

Syntax

ping mpls l2vpn (instance *instance-name* local-site-id *local-site-id-number* remote-site-id *remote-site-id-number* | interface *interface-name*)
<bottom-label-ttl>
<count *count*>
<destination *address*>
<detail>
<exp *forwarding-class*>
<logical-system (all | *logical-system-name*)>
<size *bytes*>
<source *source-address*>
<sweep>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Check the operability of MPLS Layer 2 virtual private network (VPN) connections. Type Ctrl+c to interrupt a **ping mpls l2vpn** command.

Options

bottom-label-ttl—(Optional) Display the time-to-live value for the bottom label in the label stack.

count *count*—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is **1** through 1,000,000. The default value is **5**.

destination *address*—(Optional) Specify an address other than the default (**127.0.0.1/32**) for the ping echo requests. The address can be anything within the **127/8** subnet.

detail—(Optional) Display detailed information about the echo requests sent and received.

exp *forwarding-class*—(Optional) Value of the forwarding class for the MPLS ping packets.

instance *instance-name* local-site-id *local-site-id-number* remote-site-id *remote-site-id-number*—Ping a combination of the Layer 2 VPN routing instance name, the local site identifier, and the remote site identifier, testing the integrity of the Layer 2 VPN circuit (specified by the identifiers) between the ingress and egress provider edge (PE) routers or switches.

interface *interface-name*—Ping an interface configured for the Layer 2 VPN on the egress PE router or switch.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on the specified logical system.

size *bytes*—(Optional) Size of the label-switched path (LSP) ping request packet (**96** through **65468** bytes). Packets are 4-byte aligned. For example, If you enter a size of 97, 98, 99, or 100, the router or switch uses a size value of 100 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 96-byte minimum.

source *source-address*—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (**lo.0**).

sweep—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

Additional Information

You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the egress PE router or switch (the router or switch receiving the MPLS echo packets) to ping a Layer 2 circuit.

In asymmetric MTU scenarios, the echo response may be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

Required Privilege Level

network

List of Sample Output

ping mpls l2vpn instance on page 3148
ping mpls l2vpn instance detail on page 3148

Output Fields

When you enter this command, you are provided feedback on the status of your request. An exclamation point (**!**) indicates that an echo reply was received. A period (**.**) indicates that an echo reply was not received within the timeout period. An **x** indicates that an echo reply was received with an error code these packets are not counted in the received packets count. They are accounted for separately.

ping mpls l2vpn instance

```
user@host> ping mpls l2vpn instance vpn1 remote-site-id 1 local-site-id 2
!!!!!
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

ping mpls l2vpn instance detail

```
user@host> ping mpls l2vpn instance vpn1 remote-site-id 1 local-site-id 2 detail
Request for seq 1, to interface 68, labels <800001, 100176>
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 68, labels <800001, 100176>
Reply for seq 2, return code: Egress-ok
Request for seq 3, to interface 68, labels <800001, 100176>
Reply for seq 3, return code: Egress-ok
Request for seq 4, to interface 68, labels <800001, 100176>
Reply for seq 4, return code: Egress-ok
Request for seq 5, to interface 68, labels <800001, 100176>
Reply for seq 5, return code: Egress-ok

--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

ping mpls l3vpn

Syntax

ping mpls l3vpn prefix *prefix-name*
<*l3vpn-name*>
<bottom-label-ttl>
<count *count*>
<destination *address*>
<detail>
<exp *forwarding-class*>
<logical-system (all | *logical-system-name*)>
<size *bytes*>
<source *source-address*>
<sweep>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Check the operability of a MPLS Layer 3 virtual private network (VPN) connection. Type Ctrl+c to interrupt a **ping mpls l3vpn** command.

Options

bottom-label-ttl—(Optional) Display the time-to-live value for the bottom label in the label stack.

count *count*—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is **1** through **1,000,000**. The default value is **5**.

destination *address*—(Optional) Specify an address other than the default (**127.0.0.1/32**) for the ping echo requests. The address can be anything within the **127/8** subnet.

detail—(Optional) Display detailed information about the echo requests sent and received.

exp *forwarding-class*—(Optional) Value of the forwarding class for the MPLS ping packets.

l3vpn-name—(Optional) Layer 3 VPN name.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on the specified logical system.

prefix *prefix-name*—Ping to test whether a prefix is present in a provider edge (PE) router's or switch's VPN routing and forwarding (VRF) table, by means of a Layer 3 VPN destination prefix. This option does not test the connection between a PE router or switch and a customer edge (CE) router or switch.

size *bytes*—(Optional) Size of the label-switched path (LSP) ping request packet (**96** through **65468** bytes). Packets are 4-byte aligned. For example, If you enter a size of 97, 98, 99, or 100, the router or switch uses a size value of 100 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 96-byte minimum.

source *source-address*—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (**lo.0**).

sweep—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

Additional Information

You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the egress PE router or switch (the router or switch receiving the MPLS echo packets) to ping a Layer 2 circuit.

In asymmetric MTU scenarios, the echo response may be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

Required Privilege Level

network

List of Sample Output

ping mpls l3vpn on page 3150
ping mpls l3vpn detail on page 3150

Output Fields

When you enter this command, you are provided feedback on the status of your request. An exclamation point (**!**) indicates that an echo reply was received. A period (**.**) indicates that an echo reply was not received within the timeout period. An **x** indicates that an echo reply was received with an error code these packets are not counted in the received packets count. They are accounted for separately.

ping mpls l3vpn

```
user@host> ping mpls l3vpn vpn1 prefix 10.255.245.122/32
!!!!!
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

ping mpls l3vpn detail

```
user@host> ping mpls l3vpn vpn1 prefix 10.255.245.122/32 detail
Request for seq 1, to interface 68, labels <100128, 100112>
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 68, labels <100128, 100112>
Reply for seq 2, return code: Egress-ok
Request for seq 3, to interface 68, labels <100128, 100112>
Reply for seq 3, return code: Egress-ok
Request for seq 4, to interface 68, labels <100128, 100112>
Reply for seq 4, return code: Egress-ok
Request for seq 5, to interface 68, labels <100128, 100112>
Reply for seq 5, return code: Egress-ok
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

ping mpls ldp

Syntax

ping mpls ldp *fec*
<count *count*>
<destination *address*>
<detail>
<exp *forwarding-class*>
<instance *routing-instance-name*>
<logical-system (all | *logical-system-name*)>
<size *bytes*>
<source *source-address*>
<sweep>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Check the operability of MPLS LDP-signaled label-switched path (LSP) connections. Type Ctrl+c to interrupt a **ping mpls** command.

Options

count *count*—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is **1** through **1,000,000**. The default value is **5**.

destination *address*—(Optional) Specify an address other than the default (**127.0.0.1/32**) for the ping echo requests. The address can be anything within the **127/8** subnet.

detail—(Optional) Display detailed information about the echo requests sent and received.

exp *forwarding-class*—(Optional) Value of the forwarding class for the MPLS ping packets.

fec—Ping an LDP-signaled LSP using the forwarding equivalence class (FEC) prefix and length.

instance *routing-instance-name*—(Optional) Allows you to ping a combination of the routing instance and forwarding equivalence class (FEC) associated with an LSP.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on the specified logical system.

size *bytes*—(Optional) Size of the label-switched path (LSP) ping request packet (**88** through **65468** bytes). Packets are 4-byte aligned. For example, If you enter a size of 89, 90, 91, or 92, the router or switch uses a size value of 92 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 88-byte minimum.

source *source-address*—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (**lo.0**).

sweep—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

Additional Information

If the LSP changes, the label and interface information displayed when you issued the **ping** command continues to be used. You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the remote router or switch to ping an LSP terminating there. You must configure MPLS even if you intend to ping only LDP forwarding equivalence classes (FECs).

You can configure the ping interval for the **ping mpls ldp** command by specifying a new time in seconds using the **lsp-ping-interval** statement at the **[edit protocols ldp oam]** hierarchy level. For more information, see the *Junos OS MPLS Applications Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.

In asymmetric MTU scenarios, the echo response may be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

Required Privilege Level

network

List of Sample Output

ping mpls ldp fec count on page 3152

Output Fields

When you enter this command, you are provided feedback on the status of your request. An exclamation point (**!**) indicates that an echo reply was received. A period (**.**) indicates that an echo reply was not received within the timeout period. An **x** indicates that an echo reply was received with an error code. Packets with error codes are not counted in the received packets count. They are accounted for separately.

ping mpls ldp fec count

```
user@host> ping mpls ldp 10.255.245.222 count 10
!!!xxx...x--- lsping statistics ---10 packets transmitted, 3 packets received,
70% packet loss 4 packets received with error status, not counted as received.
```

ping mpls lsp-end-point

Syntax

ping mpls lsp-end-point *prefix-name*
<count *count*>
<destination *address*>
<detail>
<exp *forwarding-class*>
<instance *routing-instance-name*>
<logical-system (all | *logical-system-name*)>
<size *bytes*>
<source *source-address*>
<sweep>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Check the operability of MPLS label-switched path (LSP) endpoint connections. Type Ctrl+c to interrupt a **ping mpls** command.

Options

count *count*—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is **1** through **1,000,000**. The default value is **5**.

destination *address*—(Optional) Specify an address other than the default (**127.0.0.1/32**) for the ping echo requests. The address can be anything within the **127/8** subnet.

detail—(Optional) Display detailed information about the echo requests sent and received.

exp *forwarding-class*—(Optional) Value of the forwarding class for the MPLS ping packets.

instance *routing-instance-name*—(Optional) Ping a combination of the routing instance and forwarding equivalence class (FEC) associated with an LSP connection.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on the specified logical system.

prefix-name—LDP forwarding equivalence class (FEC) prefix or RSVP LSP endpoint address.

size *bytes*—(Optional) Size of the LSP ping request packet. If the endpoint is LDP-based, the minimum size of the packet is **88** bytes. If the endpoint is RSVP-based, the minimum size of the packet is **100** bytes. The maximum size in either case is **65468** bytes.

source *source-address*—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (**lo.0**).

sweep—(Optional) Automatically determine the size of the maximum transmission unit (MTU).

Additional Information

If the LSP changes, the label and interface information displayed when you issued the **ping** command continues to be used. You must configure MPLS at the **[edit protocols**

mpls] hierarchy level on the remote router or switch to ping an LSP terminating there. You must configure MPLS even if you intend to ping only LDP forwarding equivalence classes (FECs).

In asymmetric MTU scenarios, the echo response may be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

Required Privilege Level network

List of Sample Output **ping mpls lsp-end-point detail on page 3154**

Output Fields When you enter this command, you are provided feedback on the status of your request. An exclamation point (**!**) indicates that an echo reply was received. A period (**.**) indicates that an echo reply was not received within the timeout period. An **x** indicates that an echo reply was received with an error code these packets are not counted in the received packets count. They are accounted for separately.

ping mpls lsp-end-point detail

```
user@host> ping mpls lsp-end-point 10.255.245.119 detail
Route to end point address is via LDP FEC
Request for seq 1, to interface 67, label 100032
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 67, label 100032
Reply for seq 2, return code: Egress-ok
Request for seq 3, to interface 67, label 100032
Reply for seq 3, return code: Egress-ok
Request for seq 4, to interface 67, label 100032
Reply for seq 4, return code: Egress-ok
Request for seq 5, to interface 67, label 100032
Reply for seq 5, return code: Egress-ok
--- lsping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
```

ping mpls rsvp

Syntax

ping mpls rsvp *lsp-name*
<count *count*>
<destination *address*>
<detail>
<dynamic-bypass>
<egress *egress-address*>
<exp *forwarding-class*>
<logical-system (all | *logical-system-name*)>
<manual-bypass>
<multipoint>
<size *bytes*>
<source *source-address*>
<standby *standby-path-name*>
<sweep>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.
The **dynamic-bypass** and **manual-bypass** options were introduced in Junos OS Release 10.2 for J-EX Series switches.

Description

Check the operability of MPLS RSVP-signaled label-switched path (LSP) connections. Type Ctrl+c to interrupt a **ping mpls** command.

Options

count *count*—(Optional) Number of ping requests to send. If **count** is not specified, five ping requests are sent. The range of values is **1** through **1,000,000**. The default value is **5**.

destination *address*—(Optional) Specify an address other than the default (**127.0.0.1/32**) for the ping echo requests. The address can be anything within the **127/8** subnet.

detail—(Optional) Display detailed information about the echo requests sent and received.

NOTE: When using the detail option, the reported time is based on the system time configured on the local and remote routers. Differences in these system times can result in an inaccurate one way ping trip times being reported.

dynamic-bypass—(Optional) Ping dynamically generated bypass LSPs, used for protecting other LSPs.

egress *egress-address*—(Optional) Only the specified egress router or switch responds to the ping request.

exp *forwarding-class*—(Optional) Value of the forwarding class for the MPLS ping packets.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on the specified logical system.

lsp-name—Ping an RSVP-signaled LSP using an LSP name.

manual-bypass—(Optional) Ping manually configured bypass LSPs, used for protecting other LSPs.

multipoint—(Optional) Send ping requests to each of the egress routers or switches participating in a point-to-multipoint LSP. You can also include the **egress** option to ping a specific egress router or switch participating in a point-to-multipoint LSP.

size *bytes*—(Optional) Size of the LSP ping request packet (**100** through **65468** bytes). Packets are 4-byte aligned. For example, if you enter a size of 101, 102, 103, or 104, the router or switch uses a size value of 104 bytes. If you enter a packet size that is smaller than the minimum size, an error message is displayed reminding you of the 100-byte minimum.

source *source-address*—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface.

standby *standby-path-name*—(Optional) Name of the standby path.

sweep —(Optional) Automatically determine the size of the maximum transmission unit (MTU).

Additional Information

If the LSP changes, the label and interface information displayed when you issued the **ping** command continues to be used. You must configure MPLS at the **[edit protocols mpls]** hierarchy level on the remote router or switch to ping an LSP terminating there. You must configure MPLS even if you intend to ping only LDP forwarding equivalence classes (FECs).

You can configure the ping interval for the **ping mpls rsvp** command by specifying a new time in seconds using the **lsp-ping-interval** statement at the **[edit protocols mpls oam]** hierarchy level. For more information, see the *Junos OS MPLS Applications Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.

In asymmetric MTU scenarios, the echo response may be dropped. For example, if the MTU from System A to System B is 1000 bytes, the MTU from System B to System A is 500 bytes, and the ping request packet size is 1000 bytes, the echo response is dropped because the PAD TLV is included in the echo response, making it too large.

Required Privilege Level

network

List of Sample Output

ping mpls rsvp (Echo Reply Received) on page 3157
ping mpls rsvp (Echo Reply with Error Code) on page 3157
ping mpls rsvp detail on page 3157
ping mpls rsvp multipoint egress detail count on page 3157
ping mpls rsvp multipoint detail count on page 3157
ping mpls rsvp destination detail count size on page 3158
ping mpls rsvp destination detail sweep size on page 3158

Output Fields

When you enter this command, you are provided feedback on the status of your request. An exclamation point (**!**) indicates that an echo reply was received. A period (**.**) indicates

that an echo reply was not received within the timeout period. An **x** indicates that an echo reply was received with an error code these packets are not counted in the received packets count. They are accounted for separately.

ping mpls rsvp (Echo Reply Received)

```
user@host> ping mpls rsvp test1
!!!!!--- lsping statistics ---5 packets transmitted, 5 packets received, 0% packet
 loss
```

ping mpls rsvp (Echo Reply with Error Code)

```
user@host> ping mpls rsvp test2
!!xxx--- lsping statistics ---5 packets transmitted, 2 packets received, 60%
packet loss3 packets received with error status, not counted as received.
```

ping mpls rsvp detail

```
user@host> ping mpls rsvp to-green detail
Request for seq 1, to interface 67, labels <100095, 0, 0>
Reply for seq 1, return code: Egress-ok
Request for seq 2, to interface 67, labels <100095, 0, 0>
Reply for seq 2, return code: Egress-ok
```

ping mpls rsvp multipoint egress detail count

```
user@host>ping mpls rsvp sample-lsp multipoint egress 192.168.1.3 detail count 1
Request for seq 1, to interface 70, label 299952
Request for seq 1, to interface 70, no label stack.
Request for seq 1, to interface 67, no label stack.

Reply for seq 1, egress 192.168.1.3, return code: Egress-ok, time: 0.242 ms
 Local transmit time: 1205310695s 215737us
 Remote receive time: 1205310695s 215979us

--- lsping, egress 192.168.1.3 statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
```

ping mpls rsvp multipoint detail count

```
user@host>ping mpls rsvp sample-lsp multipoint detail count 1
Request for seq 1, to interface 70, label 299952
Request for seq 1, to interface 70, no label stack.
Request for seq 1, to interface 67, no label stack.

Reply for seq 1, return code: Unknown TLV, time: 9.877 ms
 Local transmit time: 1205310615s 347317us
 Remote receive time: 1205310615s 357194us
Reply for seq 1, egress 192.168.1.3, return code: Egress-ok, time: 0.351 ms
 Local transmit time: 1205310615s 347262us
 Remote receive time: 1205310615s 347613us
Reply for seq 1, egress 192.168.1.13, return code: Egress-ok, time: 0.301 ms
 Local transmit time: 1205310615s 347167us
 Remote receive time: 1205310615s 347468us
Timeout for seq 1, egress 192.168.1.1
Timeout for seq 1, egress 192.168.1.4
Timeout for seq 1, egress 192.168.1.14

--- lsping, egress 192.168.1.1 statistics ---
1 packets transmitted, 0 packets received, 100% packet loss

--- lsping, egress 192.168.1.3 statistics ---
1 packets transmitted, 1 packets received, 0% packet loss

--- lsping, egress 192.168.1.4 statistics ---
1 packets transmitted, 0 packets received, 100% packet loss

--- lsping, egress 192.168.1.13 statistics ---
```

```
1 packets transmitted, 1 packets received, 0% packet loss

--- lsping, egress 192.168.1.14 statistics ---
1 packets transmitted, 0 packets received, 100% packet loss
```

ping mpls rsvp destination detail count size

```
user@host>ping mpls rsvp chaser-access destination 192.168.0.1 detail count 1 size 4468

Request for seq 1, to interface 88, label 299984, packet size 4468
Reply for seq 1, return code: Egress-ok, time: 44.804 ms
        Local transmit time: 2009-03-30 22:05:02 CEST 408.629 ms
        Remote receive time: 2009-03-30 22:05:02 CEST 453.433 ms

--- lsping statistics ---
1 packets transmitted, 1 packets received, 0% packet loss
```

ping mpls rsvp destination detail sweep size

```
user@router>  ping mpls rsvp chaser-access destination 192.168.0.1 detail sweep size 4500
Request for seq 1, to interface 86, no label stack., packet size 100
Reply for seq 1, return code: Egress-ok, time: -39.264 ms
        Local transmit time: 2009-04-24 14:05:40 CEST 541.423 ms
        Remote receive time: 2009-04-24 14:05:40 CEST 502.159 ms
Request for seq 2, to interface 86, no label stack., packet size 2300
Reply for seq 2, return code: Egress-ok, time: -38.179 ms
        Local transmit time: 2009-04-24 14:05:41 CEST 544.240 ms
        Remote receive time: 2009-04-24 14:05:41 CEST 506.061 ms
Request for seq 3, to interface 86, no label stack., packet size 4500
Timeout for seq 3
Request for seq 4, to interface 86, no label stack., packet size 3400
Reply for seq 4, return code: Egress-ok, time: -37.545 ms
        Local transmit time: 2009-04-24 14:05:45 CEST 549.953 ms
        Remote receive time: 2009-04-24 14:05:45 CEST 512.408 ms
Request for seq 5, to interface 86, no label stack., packet size 3952
Reply for seq 5, return code: Egress-ok, time: -37.176 ms
        Local transmit time: 2009-04-24 14:05:46 CEST 555.881 ms
        Remote receive time: 2009-04-24 14:05:46 CEST 518.705 ms
Request for seq 6, to interface 86, no label stack., packet size 4228
Reply for seq 6, return code: Egress-ok, time: -36.962 ms
        Local transmit time: 2009-04-24 14:05:47 CEST 561.809 ms
        Remote receive time: 2009-04-24 14:05:47 CEST 524.847 ms
Request for seq 7, to interface 86, no label stack., packet size 4368
Reply for seq 7, return code: Egress-ok, time: -36.922 ms
        Local transmit time: 2009-04-24 14:05:48 CEST 568.738 ms
        Remote receive time: 2009-04-24 14:05:48 CEST 531.816 ms
Request for seq 8, to interface 86, no label stack., packet size 4440
Reply for seq 8, return code: Egress-ok, time: -36.855 ms
        Local transmit time: 2009-04-24 14:05:49 CEST 575.669 ms
        Remote receive time: 2009-04-24 14:05:49 CEST 538.814 ms
Request for seq 9, to interface 86, no label stack., packet size 4476
Timeout for seq 9
Request for seq 10, to interface 86, no label stack., packet size 4460
Reply for seq 10, return code: Egress-ok, time: -36.906 ms
        Local transmit time: 2009-04-24 14:05:53 CEST 584.382 ms
        Remote receive time: 2009-04-24 14:05:53 CEST 547.476 ms
Request for seq 11, to interface 86, no label stack., packet size 4480
Timeout for seq 11
Request for seq 12, to interface 86, no label stack., packet size 4472
Timeout for seq 12
Request for seq 13, to interface 86, no label stack., packet size 4468
Reply for seq 13, return code: Egress-ok, time: -36.943 ms
        Local transmit time: 2009-04-24 14:06:00 CEST 594.884 ms
        Remote receive time: 2009-04-24 14:06:00 CEST 557.941 ms
Request for seq 14, to interface 86, no label stack., packet size 4476
```

```
Timeout for seq 14
Request for seq 15, to interface 86, no label stack., packet size 4472
Timeout for seq 15

--- lsp ping sweep result---
Maximum Transmission Unit (MTU) is 4468 bytes
```

request mpls lsp adjust-autobandwidth

Syntax

request mpls lsp adjust-autobandwidth
<logical-system (all | *logical-system-name*)>
<name *lsp-name*>

Syntax (J-EX Series Switch)

request mpls lsp adjust-autobandwidth
<name *lsp-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Manually trigger a bandwidth allocation adjustment for active label-switched paths (LSPs).

Options

none—Manually trigger a bandwidth allocation adjustment for all active LSP paths.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

name *lsp-name*—(Optional) Manually trigger a bandwidth allocation adjustment on the specified LSP only.

Additional Information

For this command to work properly, the following conditions must exist:

- Automatic bandwidth allocation must be enabled on the LSP. The parameters for adjustment interval and maximum average bandwidth are not reset after you issue the **request mpls lsp adjust-autobandwidth** command.
- The difference between the adjusted bandwidth and the current LSP path bandwidth must be greater than the threshold limit.

Required Privilege Level

maintenance

List of Sample Output

request mpls lsp adjust-auto-bandwidth on page 3160

Output Fields

When you enter this command, you are provided feedback on the status of your request.

request mpls lsp adjust-auto-bandwidth

```
user@host> request mpls lsp adjust-auto-bandwidth
```

show connections

Syntax

show connections
<brief | extensive>
<all | interface-switch | lsp-switch | p2mp-receive-switch | p2mp-transmit-switch | remote-interface-switch>
<down | up | up-down>
<history>
<labels>
<logical-system (all | *logical-system-name*)>
<*name*>
<status>

Syntax (J-EX Series Switch)

show connections
<brief | extensive>
<all | interface-switch | lsp-switch | p2mp-receive-switch | p2mp-transmit-switch | remote-interface-switch>
<down | up | up-down>
<history>
<labels>
<*name*>
<status>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display information about the configured circuit cross-connect (CCC) connections.

Options

none—Display the standard level of output for all configured CCC connections.

all—(Optional) Display all connections.

brief | extensive—(Optional) Display the specified level of output. Use history to display information about connection history. Use labels to display labels used for transmit and receive LSPs. Use status to display information about the connection and interface status.

interface-switch—(Optional) Display interface switch connections only.

lsp-switch—(Optional) Display LSP switch connections only.

p2mp-receive-switch—(Optional) Display point-to-multipoint LSP to local interfaces switch connections only.

p2mp-transmit-switch—(Optional) Display local interface to point-to-multipoint LSP switch connections only.

remote-interface-switch—(Optional) Display remote interface switch connections only.

down | up | up-down—(Optional) Display nonoperational, operational, or both kinds of connections.

history—(Optional) Display information about connection history.

labels—(Optional) Display labels used for transmit and receive.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

name—(Optional) Display information about the specified connection only.

status—(Optional) Display information about the connection and interface status.

Required Privilege Level view

Output Fields Table 412 on page 3162 describes the output fields for the **show connections** command. Output fields are listed in the approximate order in which they appear.

Table 412: show connections Output Fields

Field Name	Field Description
CCC and TCC connections [Link Monitoring On | Off]	Whether link monitoring is enabled: **On** or **Off**.
Legend for Status (St)	Connection or circuit status. See the output's legend for an explanation of the status field values.
Legend for connection types	Type of connection: • **if-sw**—Layer 2 switching cross-connect. • **rmt-if**—Remote interface switch. While graceful restart is in progress, **rmt-if** will display a state **(St)** of **Restart**. • **lsp-sw**—LSP stitching cross-connect. While graceful restart is in progress, **lsp-sw** will display a state **(St)** of **Restart**.
Legend for circuit types	Type of circuits: • **intf**—Interface circuit. • **tlsp**—Transmit LSP circuit. • **rlsp**—Receive LSP circuit.
Connection/Circuit	Name of the configured CCC connection.
Type	Type of connection.
St	State of the connection.
Time last up	Time that the connection or circuit last transitioned to the **Up** (operational) state.
# Up trans	Number of times that the connection or circuit has transitioned to the **Up** (operational) state.

show connections

```
user@switch> show connections
CCC and TCC connections [Link Monitoring On]
  Legend for status (St)              Legend for connection types
  UN -- uninitialized                 if-sw:  interface switching
  NP -- not present                   rmt-if: remote interface switching
  WE -- wrong encapsulation           lsp-sw: LSP switching
  DS -- disabled
  Dn -- down                          Legend for circuit types
  -> -- only outbound conn is up      intf -- interface
  <- -- only inbound  conn is up      tlsp -- transmit LSP
  Up -- operational                   rlsp -- receive LSP
  RmtDn -- remote CCC down
  Restart -- restarting

CCC Graceful restart : Restarting

Connection/Circuit                Type    St      Time last up     # Up trans
IFSW-ed                           if-sw   Up      Aug  5 15:39:15           1
  so-1/0/2.0                        intf  Up
  t1-0/1/2.0                        intf  Up
SW-db                             rmt-if  Restart                           0
  so-1/0/3.0                        intf  Up
  pro4-ca                           tlsp  Dn
  pro4-ac                           rlsp  NP
```

show connections

Syntax

show connections
<brief | extensive>
<all | remote-interface-switch>
<down | up | up-down>
<history>
<labels>
<*name*>
<status>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display information about the configured circuit cross-connect (CCC) connections.

Options

none—Display the standard level of output for all configured CCC connections on all logical systems.

brief | extensive—(Optional) Display the specified level of output.

all—(Optional) Display all connections.

down | up | up-down—(Optional) Display nonoperational, operational, or both kinds of connections.

history—(Optional) Display information about connection history.

labels—(Optional) Display labels used for transmit and receive LSPs.

name—(Optional) Display information about the specified connection only.

remote-interface-switch—(Optional) Display remote interface switch connections only.

name—(Optional) Display information about the specified connection only.

status—(Optional) Display information about the connection and interface status.

Required Privilege Level

view

Related Documentation

- Example: Configuring MPLS on J-EX Series Switches on page 3071
- Configuring MPLS on Provider Edge Switches (CLI Procedure)
- **connections on page 3128**

List of Sample Output

show connections <name> on page 3166
show connections remote-interface-switch on page 3166
show connections status on page 3167

Output Fields Table 413 on page 3165 describes the output fields for the **show connections** command. Output fields are listed in the approximate order in which they appear.

Table 413: show connections Output Fields

Field Name	Field Description
CCC and TCC connections [Link Monitoring On I Off]	Whether link monitoring is enabled: **On** or **Off**.
Legend for Status (St)	Connection or circuit status. See the output's legend for an explanation of the status field values.
Legend for connection types	Type of connection: • **if-sw**—Layer 2 switching cross-connect. • **rmt-if**—Remote interface switch. While graceful restart is in progress, **rmt-if** will display a state **(St)** of **Restart**.
Legend for circuit types	Type of circuit: • **intf**—Interface circuit. • **tlsp**—Transmit LSP circuit. • **rlsp**—Receive LSP circuit.
Connection/Circuit	Name of the configured CCC connection.
Type	Type of connection.
St	State of the connection.
Time last up	Time that the connection or circuit last transitioned to the **Up** (operational) state.
# Up trans	Number of times that the connection or circuit has transitioned to the **Up** (operational) state.

show connections

```
user@switch> show connections

Connection/Circuit               Type        St      Time last up     # Up trans
ge1-to-pe2                        rmt-if      Up      Jun 26 18:37:25
1
  ge-0/0/5.0                        intf      Up
  lsp_pe1_to_ge1_pe2                tlsp      Up
  lsp_pe2_to_ge1_pe1                rlsp      Up
```

show connections brief

```
user@switch> show connections brief
```

```
Connection/Circuit               Type        St      Time last up     # Up trans
ge-1_to_pe2                       rmt-if      Up      Jan 29 13:07:56
1
```

show connections down

```
user@switch> show connections down
No matching connections found.
```

show connections extensive

```
user@switch> show connections extensive

Connection/Circuit               Type        St      Time last up     # Up trans
ge1_to_pe2                        rmt-if      Up      Jan 29 13:07:56
1
  ge-0/0/5.0                        intf      Up
  lsp_pe1_to_ge1_pe2                tlsp      Up
  lsp_pe2_to_ge1_pe1                rlsp      Up
  Incoming labels: 299776
  Outgoing labels: Push 300112
```

show connections history

```
user@switch> show connections history

Connection/Circuit               Type        St      Time last up     # Up trans
ge1-to-pe2                        rmt-if      Up      Jan 29 13:07:56
1
                                                                     # Paths
  Time             Event                Interface/Label            Rcv Xmt
  Jan 29 13:07:56  CCC status update                                 1   1
  Jan 29 13:07:55  TLSP up              300112@1:0, 1              1   1
  Jan 29 13:07:55  TLSP down            300112@1                   1   0
  Jan 29 13:07:55  TLSP up              300112@1:0, 4097           1   1
  Jan 29 13:07:54  RLSP up              299776                     1   0
  Jan 29 13:01:08  Remote CCC down                                  0   0
  Jan 29 13:01:08  Interface up         ge-0/0/0.10                0   0
  Jan 29 13:01:06  Interface down       ge-0/0/0.10                0   0
  Jan 29 13:01:04  Remote CCC down                                  0   0
  Jan 29 13:01:02  Interface down                                   0   0
```

show connections labels

```
user@switch> show connections labels

Connection/Circuit               Type        St      Time last up     # Up trans
ge1-to-pe2                        rmt-if      RmtDn   Jun 26 18:37:25
1
  Incoming labels: 299776
  Outgoing labels: Push 299792
```

show connections <name>

```
user@switch> show connections ge1-to-pe2

Connection/Circuit               Type        St      Time last up     # Up trans
ge1_to_pe2                        rmt-if      Up      Jan 29 13:07:56
1
  ge-0/0/5.0                        intf      Up
  lsp_pe1_to_ge1_pe2                tlsp      Up
  lsp_pe2_to_ge1_pe1                rlsp      Up
```

show connections remote-interface-switch

```
Connection/Circuit               Type        St      Time last up     # Up trans
xcon10_ge0_to_239                 rmt-if      Up      Jan 29 13:07:56
       1
  ge-0/0/0.10                       intf  Up
  lsp_to_240_10                     tlsp  Up
```

```
  lsp_to_239_10                       rlsp  Up
xcon11_ge0_to_239                   rmt-if      Up      Jan 29 13:07:57
       1
  ge-0/0/0.11                         intf  Up
  lsp_to_240_11                       tlsp  Up
  lsp_to_239_11                       rlsp  Up
```

show connections status

```
user@switch> show connections status
Connection/Circuit                  Type        St      Time last up     # Up trans
xcon10_ge0_to_239                   rmt-if  Up          Jan 29 13:07:56
       1
  ge-0/0/0.10                         intf  Up
  lsp_to_240_10                       tlsp  Up
  lsp_to_239_10                       rlsp  Up
xcon11_ge0_to_239                   rmt-if  Up          Jan 29 13:07:57
       1
  ge-0/0/0.11                         intf  Up
  lsp_to_240_11                       tlsp  Up
  lsp_to_239_11                       rlsp  Up
```

show link-management

Syntax show link-management

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display Multiprotocol Label Switching (MPLS) peer and traffic engineering link information.

Options This command has no options.

Required Privilege Level view

Related Documentation

- **show link-management peer on page 3171**
- **show link-management routing on page 3173**
- **show link-management statistics on page 3176**
- **show link-management te-link on page 3178**

List of Sample Output **show link-management on page 3170**

Output Fields Table 414 on page 3168 describes the output fields for the **show link-management** command. Output fields are listed in the approximate order in which they appear.

Table 414: show link-management Output Fields

Field Name	Field Description
Peer Name	Name of the peer.
System identifier	Internal identifier for the peer. The range of values is **0** through **64,000**.
State	State of the peer: **Up** or **Down**.
Control address	Address to which a control channel is established.
CC local ID	Identifier assigned to the control channel by the local peer. The range of values is **1** through **4,294,967,296**.
CC remote ID	Identifier assigned to the control channel by the remote peer. The range of values is **1** through **4,294,967,296**.
State	State of the control channel: **Up** or **Down**.
TxSeqNum	Sequence number of the hello message being sent to the peer. The range of values is **1** through **4,294,967,295**.
RcvSeqNum	Sequence number of the last hello message received from the peer. The range of values is **0** through **4,294,967,295.**

Table 414: show link-management Output Fields *(continued)*

Field Name	Field Description
Flags	Code that provides information about the control channel. Currently supports only code value **R**, which indicates that the control channel is restarting after a failure in the control plane, as when the Link Management Protocol (LMP) process starts or restarts.
TE links	Traffic-engineered links that are managed by their peer.
TE link name	Name of the traffic-engineered link.
State	State of the traffic-engineered link: **Up**, **Down**, or **Init**.
Local identifier	Identifier of the local side of the link.
Remote identifier	Identifier of the remote side of the link.
Local address	Address of the local side of the link.
Remote address	Address of the remote side of the link.
Encoding	Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include **SDH/SONET, Ethernet**, **Packet**, and **PDH**.
Switching	Type of switching that can be performed on the traffic-engineered link. Supported values are **PSC-1** and **Packet**.
Minimum bandwidth	Smallest single allocation of bandwidth possible on the traffic-engineered link. This number is equal to the smallest bandwidth interface that is a member of the traffic-engineered link (in bps).
Maximum bandwidth	Largest single allocation of bandwidth possible on the traffic-engineered link. This number is equal to the largest bandwidth interface that is a member of the link (in bps).
Total bandwidth	Sum of the bandwidth, in bits per second (bps) and megabits per second (Mbps), of all interfaces that are members of the link.
Available bandwidth	Sum of the bandwidths of all interfaces that are members of the link and that are not yet allocated (in bps).
Name	Name of the interface.
State	State of the interface: **Up** or **Down**.
Local ID	Identifier of the local side of the interface.
Remote ID	Identifier of the remote side of the interface.
Bandwidth	Bandwidth, in bps or Mbps, of the member interface.
Used	Whether the resource is allocated to an LSP: **Yes** or **No**.

Table 414: show link-management Output Fields *(continued)*

Field Name	Field Description
LSP-name	LSP name.

show link-management

```
user@host> show link-management
Peer name: PEER-A, System identifier: 11973
 State: Up, Control address: 10.255.245.4
   CC local ID CC remote ID State       TxSeqNum  RcvSeqNum Flags
         24547        24547 Up              1027       1026
  TE links:
   pro4-ba

 TE link name: pro4-ba, State: Init
  Local identifier: 2662, Remote identifier: 0, Encoding: SDH/SONET, Switching:
  PSC-1,
  Minimum bandwidth: 155.52Mbps, Maximum bandwidth: 155.52Mbps, Total bandwidth:

  155.52Mbps,
  Available bandwidth: 155.52Mbps
    Name        State Local ID  Remote ID      Bandwidth Used  LSP-name
    so-1/0/2    Up       21271          0     155.52Mbps   No
```

show link-management peer

Syntax show link-management peer
<name *peer-name*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display Multiprotocol Label Switching (MPLS) peer link information.

Options none—Display all peer link information.

name *peer-name*—(Optional) Display information for the specified peer only.

Required Privilege Level view

Related Documentation

- **show link-management on page 3168**
- **show link-management routing on page 3173**
- **show link-management statistics on page 3176**
- **show link-management te-link on page 3178**

List of Sample Output **show link-management peer on page 3172**

Output Fields Table 415 on page 3171 describes the output fields for the **show link-management peer** command. Output fields are listed in the approximate order in which they appear.

Table 415: show link-management peer Output Fields

Field Name	Field Description
Peer Name	Name of the peer.
System identifier	Internal identifier for the peer. The range of values is **0** through **64,000**.
State	State of the peer: **Up** or **Down**.
Control address	Address to which a control channel is established.
Hello interval	How often the routing device sends Link Management Protocol (LMP) hello packets.
Hello dead interval	How long LMP waits before declaring the control channel to be dead. This is an interval during which the routing device receives no LMP hello packets from the neighbor on a control that is active or up.
CC local ID	Identifier assigned to the control channel by the local peer. The range of values is **1** through **4,294,967,296**.
CC remote ID	Identifier assigned to the control channel by the remote peer. The range of values is **1** through **4,294,967,296**.

Table 415: show link-management peer Output Fields *(continued)*

Field Name	Field Description
State	State of the control channel: **Up** or **Down**.
TxSeqNum	Sequence number of the hello message being sent to the peer. The range of values is **1** through **4,294,967,295**.
RcvSeqNum	Sequence number of the last hello message received from the peer. The range of values is **0** through **4,294,967,295.**
Flags	Code that provides information about the control channel. Currently supports only code value **R**, which indicates that the control channel is restarting after a failure in the control plane, as when the Link Management Protocol (LMP) process starts or restarts.
TE links	Traffic-engineered links that are managed by their peer.

show link-management peer

```
user@host> show link-management peer
Peer name: sonet, System identifier: 41448
  State: Up, Control address: 70.70.70.70
  Hello interval: 10000, Hello dead interval: 30000
    CC local ID CC remote ID State       TxSeqNum  RcvSeqNum Flags
           3265            0 ConfSnd            1          0 R
   TE links:
    to-sonet
```

show link-management routing

Syntax
show link-management routing
<peer <name *name*> | te-link <name *name*>>
<resource <name *name*>>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display Multiprotocol Label Switching (MPLS) peer or traffic engineering link information from the routing process.

Options
none—Display all peer and traffic-engineered link information.

peer <name *name*>—(Optional) Display information for all peers or for the specified peer only.

resource <name *name*>—(Optional) Display information for all resources or for the specified resource only.

te-link <name *name*>—(Optional) Display information for all traffic-engineered forwarding paths or for the specified path only.

Required Privilege Level
view

Related Documentation
- **show link-management on page 3168**
- **show link-management peer on page 3171**
- **show link-management statistics on page 3176**
- **show link-management te-link on page 3178**

List of Sample Output
show link-management routing on page 3175

Output Fields
Table 416 on page 3173 describes the output fields for the **show link-management routing** command. Output fields are listed in the approximate order in which they appear.

Table 416: show link-management routing Output Fields

Field Name	Field Description
Peer Name	Name of the peer.
System identifier	Internal identifier for the peer. The range of values is **0** through **64,000**.
State	State of the peer: **Up** or **Down**.
Control address	Address to which a control channel is established.
Control channel	Interface over which control packets are sent.

Table 416: show link-management routing Output Fields *(continued)*

Field Name	Field Description
State	State of the control channel.
TE link name	Traffic-engineered link name.
State	State of the traffic-engineered link: **Up** or **Down**.
Local identifier	Identifier of the local side of the link.
Remote identifier	Identifier of the remote side of the link.
Local address	Address of the local side of the link.
Remote address	Address of the remote side of the link.
Encoding	Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include **SDH/SONET**, **Ethernet**, and **Packet**.
Minimum bandwidth	Smallest single allocation of bandwidth, in bits per second (bps) or megabits per second (Mbps), possible on the traffic-engineered link. This number is equal to the smallest bandwidth interface that is a member of the traffic-engineered link.
Maximum bandwidth	Largest single allocation of bandwidth, in bps or Mbps, possible on the traffic-engineered link. This number is equal to the largest bandwidth interface that is a member of the link (in bps).
Total bandwidth	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link.
Available bandwidth	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link and that are not yet allocated.
Resource	Forwarding adjacency LSP information.
Type	Type of resource. The type is always a forwarding adjacency LSP.
State	State of the LSP: **Up** or **Down**.
System Identifier	Internal identifier for the peer. The range of values is **0** through **64,000**.
Total bandwidth	Bandwidth resource, in bps or Mbps, on the TE-link learned from the routing process.
Traffic parameters	• **Encoding**—Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include **SDH/SONET**, **Ethernet**, and **Packet**. • **Switching**—Type of switching that can be performed on the traffic-engineered link: **PSC-1** and **Packet**. • **Granularity**—Layer 2 data for switching Layer 2 LSPs for this resource. Not supported. This value is always **unknown**.

show link-management routing

```
user@host> show link-management routing
Peer name: __rpd:fe-0/1/0.0, System identifier: 2147483649
 State: Up, Control address: (null)
   Control-channel                  State
   fe-0/1/0.0                       Active

Peer name: __rpd:fe-0/1/2.0, System identifier: 2147483650
 State: Up, Control address: (null)
   Control-channel                  State
   fe-0/1/2.0                       Active

Peer name: __rpd:so-0/2/0.0, System identifier: 2147483651
 State: Down, Control address: (null)
   Control-channel                  State
   so-0/2/0.0

Peer name: __rpd:so-0/2/1.0, System identifier: 2147483652
 State: Down, Control address: (null)
   Control-channel                  State
   so-0/2/1.0

...

 TE link name: __rpd:fe-0/1/0.0, State: Up
  Local identifier: 2147483649, Remote identifier: 0,
  Local address: 192.168.37.66, Remote address: 192.168.37.66,
  Encoding: Ethernet, Minimum bandwidth: 0bps, Maximum bandwidth: 100Mbps,
  Total bandwidth: 100Mbps, Available bandwidth: 100Mbps

 TE link name: __rpd:fe-0/1/2.0, State: Up
  Local identifier: 2147483650, Remote identifier: 0,
  Local address: 192.168.37.73, Remote address: 192.168.37.73,
  Encoding: Ethernet, Minimum bandwidth: 0bps, Maximum bandwidth: 100Mbps,
  Total bandwidth: 100Mbps, Available bandwidth: 100Mbps

 TE link name: __rpd:so-0/2/0.0, State: Down
  Local identifier: 2147483651, Remote identifier: 0,
  Local address: 192.168.37.82, Remote address: 192.168.37.95,
  Encoding: Ethernet, Minimum bandwidth: 0bps, Maximum bandwidth: 155.52Mbps,
  Total bandwidth: 155.52Mbps, Available bandwidth: 155.52Mbps

...

Resource: falsp-bd, Type: LSP, State: Dn  System identifier: 2147483652,
Total bandwidth: 0bps,  Traffic parameters: Encoding: Packet, Switching: Packet,
 Granularity: Unknown

Resource: falsp-be, Type: LSP, State: Up  System identifier: 2147483654,
Total bandwidth: bw[1]=10Mbps,  Traffic parameters: Encoding: Packet,
Switching: Packet, Granularity: Unknown
```

show link-management statistics

Syntax show link-management statistics
<peer <name *name*>>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display statistical information for Link Management Protocol (LMP) packets.

Options none—Display information for all peers.

peer <name *name*>—(Optional) Display information for all peers or for the specified peer only.

Required Privilege Level view

Related Documentation

- **show link-management on page 3168**
- **show link-management peer on page 3171**
- **show link-management routing on page 3173**
- **show link-management te-link on page 3178**

List of Sample Output **show link-management statistics on page 3177**

Output Fields Table 417 on page 3176 describes the output fields for the **show link-management statistics** command. Output fields are listed in the approximate order in which they appear.

Table 417: show link-management statistics Output Fields

Field Name	Field Description
Received packets	Number of received packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
Received bad packets	Number of received bad packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
Small packets	Number of packets that are too small.
Wrong protocol version	Number of packets specifying the wrong LMP version.
Messages for unknown peer	Number of packets destined for an unknown peer.
Messages for bad state	Number of packets indicating a state that does not match the recipient.
Stale acknowledgments	Number of **configAck** and **LinkSummaryAck** packets received that have a stale message ID.
Stale negative acknowledgments	Number of **configNack** and **LinkSummaryNack** packets received that have a stale message ID.

Table 417: show link-management statistics Output Fields *(continued)*

Field Name	Field Description
Sent packets	Number of sent packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
Retransmitted packets	Number of retransmitted packets by message type. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.
Dropped packets	Number of packets sent, by message type, that have been dropped by the receiver after the LMP retransmission interval has been exceeded. If the count for a message type is zero, that message type is not displayed. If the count for all message types is zero, this field is not displayed.

show link-management statistics

```
user@host> show link-management statistics peer pro4-a
Statistics for peer pro4-a
    Received packets
        Config: 1
        Hello: 2572
    Small packets: 0
    Wrong protocol version: 0
    Messages for unknown peer: 0
    Messages for bad state: 0
    Stale acknowledgments: 0
    Stale negative acknowledgments: 0
    Sent packets
        Config: 2
        ConfigAck: 1
        Hello: 2572
    Retransmitted packets
        Config: 1
```

show link-management te-link

Syntax show link-management te-link
<brief | detail>
<name *name*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display the resources used to set up Multiprotocol Label Switching (MPLS) traffic-engineered forwarding paths.

Options none—Display information for all traffic-engineered links.

brief | detail—(Optional) Display the specified level of output.

name *name*—(Optional) Display information for the specified traffic-engineered link only.

Required Privilege Level view

Related Documentation

- **show link-management on page 3168**
- **show link-management peer on page 3171**
- **show link-management routing on page 3173**
- **show link-management statistics on page 3176**

List of Sample Output **show link-management te-link on page 3179**

Output Fields Table 418 on page 3178 describes the output fields for the **show link-management te-link** command. Output fields are listed in the approximate order in which they appear.

Table 418: show link-management te-link Output Fields

Field Name	Field Description
TE link name	Traffic-engineered link name.
State	State of the traffic-engineered link: **Up** or **Down.**
Local identifier	Identifier of the local side of the link.
Remote identifier	Identifier of the remote side of the link.
Local address	Address of the local side of the link.
Remote address	Address of the remote side of the link.
Encoding	Physical layer media type determined by the interfaces contained in the traffic-engineered link. Typical values include **SDH/SONET**, **Ethernet**, **Packet**, and **PDH**.

Table 418: show link-management te-link Output Fields *(continued)*

Field Name	Field Description
Switching	Type of switching that can be performed on the traffic-engineered link. Supported values are **PSC-1** and **Packet**.
Minimum bandwidth	Smallest single allocation of bandwidth, in bits per second (bps) or megabits per second (Mbps), possible on the traffic-engineered link. This number is equal to the smallest bandwidth interface that is a member of the traffic-engineered link.
Maximum bandwidth	Largest single allocation of bandwidth, in bps or Mbps, possible on the traffic-engineered link. This number is equal to the largest bandwidth interface that is a member of the link.
Total bandwidth	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link (in bps).
Available Bandwidth	Sum of the bandwidth, in bps or Mbps, of all interfaces that are members of the link and that are not yet allocated.
Name	Name of the interface.
State	State of the interface: **Up** or **Down**.
Local ID	Identifier of the local side of the interface.
Remote ID	Identifier of the remote side of the interface.
Bandwidth	Bandwidth, in bps or Mbps, of the member interface.
Used	Whether the resource is allocated to an LSP: **Yes** or **No**.
LSP-name	LSP name.

show link-management te-link

```
user@host> show link-management te-link
TE link name: FA-bd, State: Up
  Local identifier: 4144, Remote identifier: 0, Local address: 2.2.2.1,
  Remote address: 2.2.2.2, Encoding: Ethernet, Switching: Packet,
  Minimum bandwidth: 0bps, Maximum bandwidth: 0bps, Total bandwidth: 0bps,
Available bandwidth: 0bps
    Name        State Local ID  Remote ID      Bandwidth Used  LSP-name
    falsp-bd    Dn       43077          0           0bps No

 TE link name: FA-be, State: Up
  Local identifier: 4145, Remote identifier: 0, Local address: 1.1.1.1,
   Remote address: 1.1.1.2, Encoding: Ethernet, Switching: Packet,
  Minimum bandwidth: 0bps, Maximum bandwidth: 10Mbps, Total bandwidth: 10Mbps,
 Available bandwidth: 8Mbps
    Name        State Local ID  Remote ID      Bandwidth Used  LSP-name
    falsp-be    Up       43076          0         10Mbps  Yes  e2elsp-bf
```

show mpls admin-groups

Syntax
show mpls admin-groups
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show mpls admin-groups

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display information about configured Multiprotocol Label Switching (MPLS) administrative groups.

Options
none—Display information about the configured MPLS administrative groups.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show mpls admin-groups on page 3180

Output Fields
Table 419 on page 3180 describes the output fields for the **show mpls admin-groups** command. Output fields are listed in the approximate order in which they appear.

Table 419: show mpls admin-groups Output Fields

Field Name	Field Description
Group	Name of the administrative group.
Bit index	Value assigned to the administrative group.

show mpls admin-groups

```
user@host> show mpls admin-groups
Group        Bit index
black         3
blue          2
gold          1
green         0
```

show mpls call-admission-control

Syntax
show mpls call-admission-control
<logical-system (all | *logical-system-name*)>
<*lsp-name*>

Syntax (J-EX Series Switch)
show mpls call-admission-control
<*lsp-name*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display Multiprotocol Label Switching (MPLS) label-switched path (LSP) call admission control (CAC) information.

Options
none—Display CAC information for all LSPs.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

lsp-name—(Optional) Display CAC information for the specified LSP only.

Additional Information
The available bandwidth on an LSP path at a particular class type is the total path bandwidth at that class type minus the total bandwidth reserved by any Layer 2 connection at that class type.

Required Privilege Level
view

List of Sample Output
show mpls call-admission-control on page 3182

Output Fields
Table 420 on page 3181 describes the output fields for the **show mpls call-admission-control** command. Output fields are listed in the approximate order in which they appear.

Table 420: show mpls call-admission-control Output Fields

Field Name	Field Description
Available bandwidth	Current available bandwidth on each LSP path. Depending on whether the LSP is an E-LSP or a regular LSP, either per-class bandwidth or a single bandwidth value (corresponding to best-effort bandwidth at **ct0**) is displayed. The available bandwidth on an LSP path at a particular class type is the total path bandwidth at that class type minus the total bandwidth reserved by some Layer 2 connections at that class type.
Layer2 connections	Different Layer 2 connections that had some bandwidth requirement and were admitted into an LSP path.
LSP name	LSP pathname.
Neighbor address	Neighbor address from which CAC and bandwidth booking are configured for Layer 2 circuits.
Circuit	Interface name and circuit information.

Table 420: show mpls call-admission-control Output Fields *(continued)*

Field Name	Field Description
Primary	LSP's primary standby path.
Standby	LSP's secondary standby path.
VC bandwidth	Bandwidth constraints associated with a Layer 2 circuit route.

show mpls call-admission-control

```
user@host# show mpls call-admission-control

LSP name: pro1-be
 *Primary
  Available bandwidth: 0bps

LSP name: pro1-be-1
 *Primary
  Available bandwidth: 60kbps

LSP name: pro1-be-gold
 *Primary
  Available bandwidth: <ct0 50kbps> <ct1 20kbps> <ct2 30kbps> <ct3 0bps>
    Layer2 connections:
    Neighbor address: 10.255.245.215, Circuit: so-0/3/0.0(vc 5)
    VC bandwidth: <ct0 50kbps> <ct1 40kbps> <ct2 40kbps>

LSP name: pro1-be-gold-2
 *Primary
  Available bandwidth: <ct0 0bps> <ct1 40kbps> <ct2 40kbps> <ct3 0bps>

LSP name: pro1-be-silver
 *Primary   prim1
  Available bandwidth: <ct0 10kbps> <ct1 20kbps> <ct2 0bps> <ct3 40kbps>
    Layer2 connections:
    Neighbor address: 10.255.245.215, Circuit: so-0/3/0.1(vc 3)
    VC bandwidth: <ct0 20kbps> <ct1 20kbps>
  Standby   sec1
  Available bandwidth: <ct0 10kbps> <ct1 10kbps> <ct2 20kbps> <ct3 0bps>
    Layer2 connections:
    Neighbor address: 10.255.245.215, Circuit: so-0/3/0.1(vc 3)
    VC bandwidth: <ct0 20kbps> <ct1 20kbps>
```

show mpls cspf

Syntax
show mpls cspf
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show mpls cspf

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display Multiprotocol Label Switching (MPLS) Constrained Shortest Path First (CSPF) statistics.

Options
none—Display MPLS CSFP statistics.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show mpls cspf on page 3184

Output Fields
Table 421 on page 3183 describes the output fields for the **show mpls cspf** command. Output fields are listed in the approximate order in which they appear.

Table 421: show mpls cspf Output Fields

Field Name	Field Description
Queue length	Number of LSPs queued for automatic path computation.
current	Current queue length.
maximum	Maximum queue length (high-water mark).
dequeued	Number of aborted computation attempts.
Paths	Counters for label-switched path computations.
total	Sum of the next four fields.
successful	Number of path computations that were successfully completed.
no route	Number of path computations that failed because the destination is unreachable.
Sys Error	Number of path computations that failed because of lack of memory.
CSPFs	Total number of CSPF computations. A single path might require multiple CSPF computations.

Table 421: show mpls cspf Output Fields *(continued)*

Field Name	Field Description
Time	Time, in seconds, required to perform the label-switched path computation.
Total	Total amount of time consumed by the CSPF path computation algorithm.
CSPFs	Total number of CSPF computations.
Avg per CSPF	Average amount of time required for each CSPF computation.
% of rpd	Percentage of routing process CPU used in the CSPF computation.

show mpls cspf

```
user@host> show mpls cspf
CSPF statistics
Queue length   current       maximum      dequeued
                     0             0             0
Paths            total    successful     no route     sys error         CSPFs
                     0             0             0             0             0
Time (secs)      total         CSPFs  avg per CSPF      % of rpd
                0.000000      0.000000      0.000000        0.0000
```

show mpls diffserv-te

Syntax show mpls diffserve-te
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch) show mpls diffserve-te

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display Multiprotocol Label Switching (MPLS) label-switched path (LSP) Differentiated Services (DiffServ) class and preemption priority information.

Options none—Display DiffServ classes and priorities used by MPLS LSPs.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level view

List of Sample Output **show mpls diffserv-te on page 3185**

Output Fields Table 422 on page 3185 describes the output fields for the **show mpls diffserv-te** command. Output fields are listed in the approximate order in which they appear.

Table 422: show mpls diffserv-te Output Fields

Field Name	Field Description
Bandwidth model	Bandwidth constraint model supported. The maximum allocation model (MAM) for EXP-inferred LSPs (E-LSPs) is currently supported.
TE class	DiffServ traffic engineering class.
Traffic class	MPLS class type that corresponds to the DiffServ traffic engineering class: • **ct0**—Best effort • **ct1**—Assured forwarding • **ct2**—Expedited forwarding • **ct3**—Network control
Priority	MPLS preemption priority for this class type, a value from **0** through **7**. Interior gateway protocols (IGPs) distribute information about the available bandwidth for each traffic engineering class.

show mpls diffserv-te

```
user@host> show mpls diffserv-te
Bandwidth model: Maximum Allocation Model with support for E-LSPs.
TE class     Traffic class       Priority
```

```
te0             ct0                 3
te1             ct1                 2
```

show mpls interface

Syntax
show mpls interface
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show mpls interface

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display information about Multiprotocol Label Switching (MPLS)-enabled interfaces.

Options
none—Display information about MPLS-enabled interfaces.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Additional Information
MPLS is enabled on an interface when the interface is configured with both the **set protocol mpls interface** ***interface-name*** and **set interface** ***interface-name*** **unit 0 family mpls** statements.

Required Privilege Level
view

List of Sample Output
show mpls interface on page 3187

Output Fields
Table 423 on page 3187 describes the output fields for the **show mpls interface** command. Output fields are listed in the approximate order in which they appear.

Table 423: show mpls interface Output Fields

Field Name	Field Description
Interface	Name of the interface.
State	State of the interface: **Up** or **Dn** (down).
Administrative groups	Administratively assigned colors of the link.

show mpls interface

```
user@host> show mpls interface
Interface   State       Administrative groups
so-1/0/0.0  Up          Blue Yellow Red
```

show mpls interface

Syntax show mpls interface

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display information about MPLS-enabled interfaces. MPLS is enabled on an interface when the interface is configured with both the **set protocols mpls interface** ***interface-name*** and **set interfaces** ***interface-name*** **unit 0 family mpls** commands.

Required Privilege Level view

Related Documentation

- Example: Configuring MPLS on J-EX Series Switches on page 3071
- Configuring MPLS on Provider Edge Switches (CLI Procedure)
- Configuring MPLS on Provider Switches (CLI Procedure) on page 3102

List of Sample Output **show mpls interface on page 3188**

Output Fields Table 424 on page 3188 describes the output fields for the **show mpls interface** command. Output fields are listed in the approximate order in which they appear.

Table 424: show mpls interface Output Fields

Field Name	Field Description
Interface	Name of the interface.
State	State of the interface: **Up** or **Dn** (down).
Administrative groups	Administratively assigned colors of the link.

show mpls interface

```
user@switch> show mpls interface
Interface   State       Administrative groups
so-1/0/0.0  Up          Blue Yellow Red
```

show mpls lsp

Syntax

show mpls lsp
<brief | detail | extensive | terse>
<bidirectional | unidirectional>
<bypass>
<defaults>
<descriptions>
<down | up>
<logical-system (all | *logical-system-name*)>
<*lsp-type*>
<name *name*>
<p2mp>
<statistics>
<transit>

Syntax (J-EX Series Switch)

show mpls lsp
<brief | detail | extensive | terse>
<bidirectional | unidirectional>
<bypass>
<descriptions>
<down | up>
<*lsp-type*>
<name *name*>
<p2mp>
<statistics>
<transit>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display information about configured and active dynamic Multiprotocol Label Switching (MPLS) label-switched paths (LSPs).

Options

none—Display standard information about all configured and active dynamic MPLS LSPs.

brief | detail | extensive | terse—(Optional) Display the specified level of output. The extensive option displays the same information as the detail option, but covers the most recent 50 events.

bidirectional | unidirectional—(Optional) Display bidirectional or unidirectional LSP information, respectively.

bypass—(Optional) Display LSPs used for protecting other LSPs.

defaults—(Optional) Display the MPLS LSP default settings.

descriptions—(Optional) Display the MPLS label-switched path (LSP) descriptions. To view this information, you must configure the description statement at the **[edit protocol mpls lsp]** hierarchy level. Only LSPs with a description are displayed. This command is only valid for the ingress routing device, because the description is not propagated in RSVP messages.

down | up—(Optional) Display only LSPs that are inactive or active, respectively.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

lsp-type—(Optional) Display information about a particular LSP type:

- **bypass**—Sessions for bypass LSPs.
- **egress**—Sessions that terminate on this routing device.
- **ingress**—Sessions that originate from this routing device.
- **transit**—Sessions that pass through this routing device.

name *name*—(Optional) Display information about the specified LSP or group of LSPs.

p2mp—(Optional) Display information about point-to-multipoint LSPs.

statistics—(Optional) (Egress and transit routers only) Display accounting information about LSPs. Statistics are not available for LSPs on the egress routing device, because the penultimate routing device in the LSP sets the label to 0. Also, as the packet arrives at the egress routing device, the hardware removes its MPLS header and the packet reverts to being an IPv4 packet. Therefore, it is counted as an IPv4 packet, not an MPLS packet.

transit—(Optional) Display LSPs transiting this routing device.

Required Privilege Level view

Related Documentation
- **clear mpls lsp on page 3140**

List of Sample Output
show mpls lsp defaults on page 3195
show mpls lsp descriptions on page 3196
show mpls lsp detail on page 3196
show mpls lsp extensive on page 3196
show mpls lsp p2mp on page 3197
show mpls lsp p2mp detail on page 3197

Output Fields Table 425 on page 3190 describes the output fields for the **show mpls lsp** command. Output fields are listed in the approximate order in which they appear.

Table 425: show mpls lsp Output Fields

Field Name	Field Description	Level of Output
Ingress LSP	Information about LSPs on the ingress routing device. Each session has one line of output.	All levels
Egress LSP	Information about the LSPs on the egress routing device. MPLS learns this information by querying RSVP, which holds all the transit and egress session information. Each session has one line of output.	All levels

Table 425: show mpls lsp Output Fields *(continued)*

Field Name	Field Description	Level of Output
Transit LSP	Number of LSPs on the transit routing devices and the state of these paths. MPLS learns this information by querying RSVP, which holds all the transit and egress session information.	All levels
P2MP name	Name of the point-to-multipoint LSP. Dynamically generated P2MP LSPs used for VPLS flooding use dynamically generated P2MP LSP names. The name uses the format ***identifier*:vpls:*router-id*:*routing-instance-name***. The ***identifier*** is automatically generated by the Junos OS.	All levels
P2MP branch count	Number of destination LSPs the point-to-multipoint LSP is transmitting to.	All levels
P	An asterisk (*) under this heading indicates that the LSP is a primary path.	All levels
address	(**detail** and **extensive**) Destination (egress routing device) of the LSP.	**detail extensive**
To	Destination (egress routing device) of the session.	**brief**
From	Source (ingress routing device) of the session.	**brief detail**
State	State of the LSP handled by this RSVP session: **Up**, **Dn** (down), or **Restart**.	**brief detail**
Active Route	Number of active routes (prefixes) installed in the forwarding table. For ingress LSPs, the forwarding table is the primary IPv4 table (**inet.0**). For transit and egress RSVP sessions, the forwarding table is the primary MPLS table (**mpls.0**).	**detail extensive**
P	Path. An asterisk (*) underneath this column indicates that the LSP is a primary path.	**brief**
LSPname	Name of the LSP.	**brief detail**
DiffServeInfo	Type of LSP: multiclass LSP **(multiclass diffServ-TE LSP)** or Differentiated-Services-aware traffic engineering LSP **(diffServ-TE LSP)**.	**detail**
Bypass	(Bypass LSP) Destination address (egress routing device) for the bypass LSP.	All levels
LSPpath	Indicates whether the RSVP session is for the primary or secondary LSP path. **LSPpath** can be either **primary** or **secondary** and can be displayed on the ingress, egress, and transit routing devices.	**detail**
Bidir	(GMPLS) The LSP allows data to travel in both directions between GMPLS devices.	All levels
Bidirectional	(GMPLS) The LSP allows data to travel both ways between GMPLS devices.	All levels
Rt	Number of active routes (prefixes) installed in the routing table. For ingress RSVP sessions, the routing table is the primary IPv4 table (**inet.0**). For transit and egress RSVP sessions, the routing table is the primary MPLS table (mpls.0).	**brief**
ActivePath	(Ingress LSP) Name of the active path: **Primary** or **Secondary**.	**detail extensive**

Table 425: show mpls lsp Output Fields *(continued)*

Field Name	Field Description	Level of Output
FastReroute desired	Fast reroute has been requested by the ingress routing device.	**detail**
Link protection desired	Link protection has been requested by the ingress routing device.	**detail**
LoadBalance	(Ingress LSP) CSPF load-balancing rule that was configured to select the LSP's path among equal-cost paths: **Most-fill**, **Least-fill**, or **Random**.	**detail extensive**
Signal type	Signal type for GMPLS LSPs. The signal type determines the peak data rate for the LSP: **DS0**, **DS3**, **STS-1**, **STM-1**, or **STM-4**.	All levels
Encoding type	LSP encoding type: **Packet**, **Ethernet**, **PDH**, **SDH/SONET**, **Lambda**, or **Fiber**.	All levels
Switching type	Type of switching on the links needed for the LSP: **Fiber**, **Lamda**, **Packet**, **TDM**, or **PSC-1**.	All levels
GPID	Generalized Payload Identifier (identifier of the payload carried by an LSP): **HDLC**, **Ethernet**, **IPv4**, **PPP**, or **Unknown**.	All levels
Protection	Configured protection capability desired for the LSP: **Extra**, **Enhanced**, **none**, **One plus one**, **One to one**, or **Shared**.	All levels
Upstream label in	(Bidirectional LSPs) Incoming label for reverse direction traffic for this LSP.	All levels
Upstream label out	(Bidirectional LSPs) Outgoing label for reverse direction traffic for this LSP.	All levels
Suggested label received	(Bidirectional LSPs) Label the upstream node suggests to use in the Resv message that is sent.	All levels
Suggested label sent	(Bidirectional LSPs) Label the downstream node suggests to use in the Resv message that is returned.	All levels
Autobandwidth	(Ingress LSP) The LSP is performing autobandwidth allocation.	**detail extensive**
MinBW	(Ingress LSP) Configured minimum value of the LSP, in bps.	**detail extensive**
MaxBW	(Ingress LSP) Configured maximum value of the LSP, in bps.	**detail extensive**
AdjustTimer	(Ingress LSP) Configured value of the bandwidth adjustment timer, indicating the total amount of time allowed before bandwidth adjustment will take place, in seconds.	**detail extensive**
MaxAvgBW util	(Ingress LSP) Current value of the actual maximum average bandwidth utilization, in bps.	**detail extensive**
Overflow limit	(Ingress LSP) Configured value of the threshold overflow limit.	**detail extensive**

Table 425: show mpls lsp Output Fields *(continued)*

Field Name	Field Description	Level of Output
Overflow sample count	(Ingress LSP) Current value for the overflow sample count.	**detail extensive**
Bandwidth Adjustment in *nnn* second(s)	(Ingress LSP) Current value of the bandwidth adjustment timer, indicating the amount of time remaining until the bandwidth adjustment will take place, in seconds.	**detail extensive**
Active path indicator	(Ingress LSP) A value of * indicates that the path is active. The absence of * indicates that the path is not active. In the following example, "long" is the active path. ***Primary long** **Standby short**	**detail extensive**
Primary	(Ingress LSP) Name of the primary path.	**detail extensive**
Secondary	(Ingress LSP) Name of the secondary path.	**detail extensive**
Standby	(Ingress LSP) Name of the path in standby mode.	**detail extensive**
State	(Ingress LSP) State of the path: **Up** or **Dn** (down).	**detail extensive**
COS	(Ingress LSP) Class-of-service value.	**detail extensive**
Bandwidth per class	(Ingress LSP) Active bandwidth for the LSP path for each MPLS class type, in bps.	**detail extensive**
OptimizeTimer	(Ingress LSP) Configured value of the optimize timer, indicating the total amount of time allowed before path reoptimization, in seconds.	**detail extensive**
SmartOptimizeTimer	(Ingress LSP) Configured value of the smart optimize timer, indicating the total amount of time allowed before path reoptimization, in seconds.	**detail extensive**
Reoptimization in xxx seconds	(Ingress LSP) Current value of the optimize timer, indicating the amount of time remaining until the path will be reoptimized, in seconds.	**detail extensive**
Computed ERO (S [L] denotes strict [loose] hops)	(Ingress LSP) Computed explicit route. A series of hops, each with an address followed by a hop indicator. The value of the hop indicator can be strict (**S**) or loose (**L**).	**detail extensive**
CSPF metric	(Ingress LSP) Constrained Shortest Path First metric for this path.	**detail extensive**

Table 425: show mpls lsp Output Fields *(continued)*

Field Name	Field Description	Level of Output
Received RRO	(Ingress LSP) Received record route. A series of hops, each with an address followed by a flag. (In most cases, the received record route is the same as the computed explicit route. If **Received RRO** is different from **Computed ERO**, there is a topology change in the network, and the route is taking a detour.) The following flags identify the protection capability and status of the downstream node: • **0x01**—Local protection available. The link downstream from this node is protected by a local repair mechanism. This flag can be set only if the Local protection flag was set in the **SESSION_ATTRIBUTE** object of the corresponding Path message. • **0x02**—Local protection in use. A local repair mechanism is in use to maintain this tunnel (usually because of an outage of the link it was routed over previously). • **0x03**—Combination of **0x01** and **0x02**. • **0x04**—Bandwidth protection. The downstream routing device has a backup path providing the same bandwidth guarantee as the protected LSP for the protected section. • **0x08**—Node protection. The downstream routing device has a backup path providing protection against link and node failure on the corresponding path section. If the downstream routing device can set up only a link-protection backup path, the **Local protection available** bit is set but the **Node protection** bit is cleared. • **0x09**—Detour is established. Combination of **0x01** and **0x08**. • **0x10**—Preemption pending. The preempting node sets this flag if a pending preemption is in progress for the traffic engine LSP. This flag indicates to the ingress legacy edge router (LER) of this LSP that it should be rerouted. • **0xb**—Detour is in use. Combination of **0x01**, **0x02**, and **0x08**.	**detail extensive**
Index number	(Ingress LSP) Log entry number of each LSP path event. The numbers are in chronological descending order, with a maximum of 50 index numbers displayed.	**extensive**
Date	(Ingress LSP) Date of the LSP event.	**extensive**
Time	(Ingress LSP) Time of the LSP event.	**extensive**
Event	(Ingress LSP) Description of the LSP event.	**extensive**
Created	(Ingress LSP) Date and time the LSP was created.	**extensive**
Resv style	(Bypass) RSVP reservation style. This field consists of two parts. The first is the number of active reservations. The second is the reservation style, which can be **FF** (fixed filter), **SE** (shared explicit), or **WF** (wildcard filter).	**brief detail extensive**
Labelin	Incoming label for this LSP.	**brief detail**
Labelout	Outgoing label for this LSP.	**brief detail**
LSPname	Name of the LSP.	**brief detail**

Table 425: show mpls lsp Output Fields *(continued)*

Field Name	Field Description	Level of Output
Time left	Number of seconds remaining in the lifetime of the reservation.	**detail**
Since	Date and time when the RSVP session was initiated.	**detail**
Tspec	Sender's traffic specification, which describes the sender's traffic parameters.	**detail**
Port number	Protocol ID and sender or receiver port used in this RSVP session.	**detail**
PATH rcvfrom	Address of the previous-hop (upstream) routing device or client, interface the neighbor used to reach this router, and number of packets received from the upstream neighbor.	**detail**
PATH sentto	Address of the next-hop (downstream) routing device or client, interface used to reach this neighbor, and number of packets sent to the downstream routing device.	**detail**
RESV rcvfrom	Address of the previous-hop (upstream) routing device or client, interface the neighbor used to reach this routing device, and number of packets received from the upstream neighbor. The output in this field, which is consistent with that in the **PATH rcvfrom** field, indicates that the RSVP negotiation is complete.	**detail**
Record route	Recorded route for the session, taken from the record route object.	**detail**
Soft preempt	Number of soft preemptions that occurred on a path and when the last soft preemption occurred. Only successful soft preemptions are counted (those that actually resulted in a new path being used).	**detail**
Soft preemption pending	Path is in the process of being soft preempted. This display is removed once the ingress router has calculated a new path.	**detail**
MPLS-TE LSP Defaults	Default settings for MPLS traffic engineered LSPs: • **LSP Holding Priority**—Determines the degree to which an LSP holds on to its session reservation after the LSP has been set up successfully. • **LSP Setup Priority**—Determines whether a new LSP that preempts an existing LSP can be established. • **Hop Limit**—Specifies the maximum number of routers the LSP can traverse (including the ingress and egress). • **Bandwidth**—Specifies the bandwidth in bits per second for the LSP. • **LSP Retry Timer**—Length of time in seconds that the ingress router waits between attempts to establish the primary path.	**defaults**

show mpls lsp defaults

```
user@host> show mpls lsp defaults
   MPLS-TE LSP Defaults
      LSP Holding Priority      0
      LSP Setup Priority        7
      Hop Limit                 255
```

```
    Bandwidth                0
    LSP Retry Timer          30 seconds
```

show mpls lsp descriptions

```
user@host> show mpls lsp descriptions
Ingress LSP: 3 sessions
To              LSP name                                Description
10.0.0.195      to-sanjose                              to-sanjose-desc
10.0.0.195      to-sanjose-other-desc                   other-desc
Total 2 displayed, Up 2, Down 0
```

show mpls lsp detail

```
user@host> show mpls lsp detail
Ingress LSP: 1 sessions

10.255.245.3
  From: 10.255.245.5, State: Up, ActiveRoute: 1, LSPname: lsp-ec
  ActivePath: long-path (primary)
  LoadBalance: Random
  Autobandwidth
  MaxBW: 5Mbps
  AdjustTimer: 4800 secs AdjustThreshold: 1%
  Max AvgBW util: 0bps, Bandwidth Adjustment in 3383 second(s).
  Overflow limit: 5, Overflow sample count: 0
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   long-path        State: Up
    SmartOptimizeTimer: 180
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 5)
 192.168.37.89 S 192.168.37.87 S
   Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

          192.168.37.89 192.168.37.87
Total 1 displayed, Up 1, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

show mpls lsp extensive

```
user@host> show mpls lsp extensive
Ingress LSP: 5 sessions

10.255.71.242
  From: 10.255.71.238, State: Up, ActiveRoute: 1009, LSPname: sample-ccc
  ActivePath: path3 (primary)
  Link protection desired
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   path3            State: Up
    OptimizeTimer: 30
    SmartOptimizeTimer: 180
    Reoptimization in 26 second(s).
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 1)
 10.35.1.41 S
   Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

          10.35.1.41(Label=3)
   10 Dec  8 13:51:58.986 CSPF: computation result ignored
    9 Dec  8 13:51:30.547 Record Route:  10.35.1.41(Label=3)
    8 Dec  8 13:51:30.547 Up
    7 Dec  8 13:51:30.397 Originate make-before-break call
    6 Dec  8 13:51:30.397 CSPF: computation result accepted  10.35.1.41
    5 Dec  8 13:50:41.467 Selected as active path
    4 Dec  8 13:50:41.467 Record Route:  10.35.1.41(Label=3)
```

```
    3 Dec  8 13:50:41.466 Up
    2 Dec  8 13:50:41.371 Originate Call
    1 Dec  8 13:50:41.371 CSPF: computation result accepted  10.35.1.41
  Created: Fri Dec  8 13:50:40 2006
Total 1 displayed, Up 1, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

show mpls lsp p2mp

```
user@host> show mpls lsp p2mp
Ingress LSP: 2 sessions
P2MP name: p2mp-lsp1, P2MP branch count: 1
To              From            State Rt ActivePath       P     LSPname
10.255.245.51   10.255.245.50   Up     0 path1            *     p2mp-branch-1
P2MP name: p2mp-lsp2, P2MP branch count: 1
To              From            State Rt ActivePath       P     LSPname
10.255.245.51   10.255.245.50   Up     0 path1            *     p2mp-st-br1
Total 2 displayed, Up 2, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

show mpls lsp p2mp detail

```
user@host> show mpls lsp p2mp detail
Ingress LSP: 2 sessions
P2MP name: p2mp-lsp1, P2MP branch count: 1

10.255.245.51
  From: 10.255.245.50, State: Up, ActiveRoute: 0, LSPname: p2mp-branch-1
  ActivePath: path1 (primary)
  P2MP name: p2mp-lsp1
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   path1            State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
 192.168.208.17 S
   Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

          192.168.208.17
P2MP name: p2mp-lsp2, P2MP branch count: 1

10.255.245.51
  From: 10.255.245.50, State: Up, ActiveRoute: 0, LSPname: p2mp-st-br1
  ActivePath: path1 (primary)
  P2MP name: p2mp-lsp2
  LoadBalance: Random
  Encoding type: Packet, Switching type: Packet, GPID: IPv4
 *Primary   path1            State: Up
    Computed ERO (S [L] denotes strict [loose] hops): (CSPF metric: 25)
 192.168.208.17 S
   Received RRO (ProtectionFlag 1=Available 2=InUse 4=B/W 8=Node 10=SoftPreempt):

          192.168.208.17
Total 2 displayed, Up 2, Down 0
```

show mpls path

Syntax

show mpls path
<logical-system (all | *logical-system-name*)>
<*path-name*>

Syntax (J-EX Series Switch)

show mpls path
<*path-name*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display dynamic Multiprotocol Label Switching (MPLS) label-switched paths (LSPs).

Options

none—Display standard information about all MPLS LSPs.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

path-name—(Optional) Display information about the specified LSP only.

Required Privilege Level

view

List of Sample Output

show mpls path on page 3198

Output Fields

Table 426 on page 3198 describes the output fields for the **show mpls path** command. Output fields are listed in the approximate order in which they appear.

Table 426: show mpls path Output Fields

Field Name	Field Description
Path name	Information about ingress LSPs. Each path has one line of output.
Address	Addresses of the routing devices that form the LSP.
Strict/loose address	Whether the address is a configured as a strict or loose address.

show mpls path

```
user@host> show mpls path
Path name        Address          Strict/loose address
p1               123.456.55.6     Strict
                 123.456.1.6      Loose
p2               191.456.1.4      Strict
```

show route forwarding-table

Syntax

show route forwarding-table
<detail | extensive | summary>
<ccc *ccc-interface-name*>
<*destination*>
<family *family-name*>
<label *label*>
<matching *ip_prefix*>
<multicast>
<vpn *vpn*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display the Routing Engine's forwarding table, including the network-layer prefixes and their next hops. This command is used to help verify that the routing protocol process has relayed the correction information to the forwarding table. The Routing Engine constructs and maintains one or more routing tables. From the routing tables, the Routing Engine derives a table of active routes, called the forwarding table.

Options

none—Display the routes in the forwarding table.

detail | extensive | summary—(Optional) Display the specified level of output.

ccc—(Optional) Display the specified circuit cross-connect interface name for entries to match.

destination —(Optional) Display the destination prefix.

family *family-name* —(Optional) Display routing table entries for the specified family: **ethernet-switching**, **inet**, **inet6**, **iso**, **mpls**, **vlan classification**.

label *label* —(Optional) Display route entries for the specified label name.

matching *ip_prefix* —(Optional) Display route entries for the specified IP prefix.

multicast—(Optional) Display route entries for multicast routes.

vpn *vpn* —(Optional) Display route entries for the specified VPN.

Required Privilege Level

view

Related Documentation

- Example: Configuring MPLS on J-EX Series Switches on page 3071
- Configuring MPLS on Provider Edge Switches (CLI Procedure)
- Configuring MPLS on Provider Switches (CLI Procedure) on page 3102

List of Sample Output

show route forwarding-table on page 3201
show route forwarding-table summary on page 3202
show route forwarding-table extensive on page 3202

show route forwarding-table ccc on page 3203
show route forwarding-table family on page 3204
show route forwarding-table label on page 3204
show route forwarding-table matching on page 3204
show route forwarding-table multicast on page 3205

Output Fields Table 427 on page 3200 lists the output fields for the **show route forwarding-table** command. Output fields are listed in the approximate order in which they appear. Field names might be abbreviated (as shown in parentheses) when no level of output is specified or when the **detail** keyword is used instead of the **extensive** keyword.

Table 427: show route forwarding-table Output Fields

Field Name	Field Description	Level of Output
Routing table	Name of the routing table (for example, **inet**, **inet6**, **mpls**).	All levels
Address family	Address family (for example, **IP**, **IPv6**, **ISO**, **MPLS**).	All levels
Destination	Destination of the route.	**detail, extensive**
Route Type (Type)	How the route was placed into the forwarding table. When the **detail** keyword is used, the route type might be abbreviated (as shown in parentheses): • **cloned** (**clon**)—(TCP or multicast only) Cloned route. • **destination** (**dest**)—Remote addresses directly reachable through an interface. • **destination down** (**iddn**)—Destination route for which the interface is unreachable. • **interface cloned** (**ifcl**)—Cloned route for which the interface is unreachable. • **route down** (**ifdn**)—Interface route for which the interface is unreachable. • **ignore** (**ignr**)—Ignore this route. • **interface** (**intf**)—Installed as a result of configuring an interface. • **permanent** (**perm**)—Routes installed by the kernel when the routing table is initialized. • **user**—Routes installed by the routing protocol process or as a result of the configuration.	All levels
Route reference (RtRef)	Number of routes to reference.	**detail, extensive**
Flags	Route type flags: • **none**—No flags are enabled. • **accounting**—Route has accounting enabled. • **cached**—Cache route. • **incoming-iface** ***interface-number*** —Check against incoming interface. • **prefix load balance**—Load balancing is enabled for this prefix. • **sent to PFE**—Route has been sent to the Packet Forwarding Engine. • **static**—Static route.	**extensive**
Nexthop	IP address of the next hop to the destination.	**detail, extensive**

Table 427: show route forwarding-table Output Fields *(continued)*

Field Name	Field Description	Level of Output
Next hop type (Type)	Next-hop type. When the **detail** keyword is used, the next-hop type might be abbreviated (as indicated in parentheses): • **broadcast (bcst)**—Broadcast. • **deny**—Deny. • **hold**—Next hop is waiting to be resolved into a unicast or multicast type. • **indexed (idxd)**—Indexed next hop. • **indirect (indr)**—Indirect next hop. • **local (locl)**—Local address on an interface. • **routed multicast (mcrt)**—Regular multicast next hop • **multicast (mcst)**—Wire multicast next hop (limited to the LAN). • **multicast discard (mdsc)**—Multicast discard. • **multicast group (mgrp)** —Multicast group member. • **receive (recv)**—Receive. • **reject (rjct)** Discard. An ICMP unreachable message was sent. • **resolve (rslv)**—Resolving the next hop. • **unicast (ucst)**—Unicast. • **unilist (ulst)**—List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.	**detail, extensive**
Index	Software index of the next hop that is used to route the traffic for a given prefix.	**detail, extensive** none
Route interface-index	Logical interface index from which the route is learned. For example, for interface routes, this is the logical interface index of the route itself. For static routes, this field is zero. For routes learned through routing protocols, this is the logical interface index from which the route is learned.	**extensive**
Reference (NhRef)	Number of routes that refer to this next hop.	none **detail, extensive**
Next-hop interface (Netif)	Interface used to reach the next hop.	none **detail, extensive**
Alternate forward nh index	Index number of the alternate next hop interface. Seen with **multicast** option only.	**extensive**
Next-hop L3 Interface	The next hop layer 3 interface. This option can be expressed as a VLAN name and is only seen with the **multicast** option.	**extensive**
Next-hop L2 Interfaces	The next hop layer 2 interfaces. Seen with **multicast** option only.	**extensive**

show route forwarding-table

```
user@switch> show route forwarding-table

Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            user     2 0:12:f2:21:cf:0    ucst   333     5 me0.0
default            perm     0                    rjct    36     2
```

```
0.0.0.0/32          perm     0                    dscd    34     1
2.2.2.0/24          intf     0                    rslv  1309     1 ae0.0
2.2.2.0/32          dest     0 2.2.2.0            recv  1307     1 ae0.0
2.2.2.1/32          dest     0 0:21:59:cc:89:c0   ucst  1320     1 ae0.0
2.2.2.2/32          intf     0 2.2.2.2            locl  1308     2
2.2.2.2/32          dest     0 2.2.2.2            locl  1308     2
2.2.2.255/32        dest     0 2.2.2.255          bcst  1306     1 ae0.0
3.3.3.0/24          intf     0                    rslv  1313     1 ae1.0
3.3.3.0/32          dest     0 3.3.3.0            recv  1311     1 ae1.0
3.3.3.1/32          intf     0 3.3.3.1            locl  1312     2
3.3.3.1/32          dest     0 3.3.3.1            locl  1312     2
3.3.3.2/32          dest     0 0:21:59:cc:89:c1   ucst  1321    24 ae1.0
3.3.3.255/32        dest     0 3.3.3.255          bcst  1310     1 ae1.0
4.4.4.0/24          user     0 3.3.3.2            ucst  1321    24 ae1.0
8.8.8.8/32          user     0 3.3.3.2            ucst  1321    24 ae1.0
9.9.9.9/32          intf     0 9.9.9.9            locl  1280     1
10.10.10.10/32      user     0 3.3.3.2            ucst  1321    24 ae1.0
10.93.8.0/21        intf     0                    rslv   323     1 me0.0
10.93.8.0/32        dest     0 10.93.8.0          recv   321     1 me0.0
10.93.13.238/32     intf     0 10.93.13.238       locl   322     2
10.93.13.238/32     dest     0 10.93.13.238       locl   322     2
10.93.15.254/32     dest     0 0:12:f2:21:cf:0    ucst   333     5 me0.0
10.93.15.255/32     dest     0 10.93.15.255       bcst   320     1 me0.0
14.14.14.0/24       ifdn     0                    rslv  1319     1 ge-0/0/25.0
14.14.14.0/32       iddn     0 14.14.14.0         recv  1317     1 ge-0/0/25.0
14.14.14.2/32       user     0                    rjct    36     2
14.14.14.2/32       intf     0 14.14.14.2         locl  1318     2
14.14.14.2/32       iddn     0 14.14.14.2         locl  1318     2
14.14.14.255/32     iddn     0 14.14.14.255       bcst  1316     1 ge-0/0/25.0
224.0.0.0/4         perm     1                    mdsc    35     1
224.0.0.1/32        perm     0 224.0.0.1          mcst    31     3
224.0.0.5/32        user     1 224.0.0.5          mcst    31     3
255.255.255.255/32 perm      0                    bcst    32     1
```

show route forwarding-table summary

```
user@switch> show route forwarding-table summary

Routing table: default.inet
Internet:
         user:          6 routes
         perm:          5 routes
         intf:          8 routes
         dest:         12 routes
         ifdn:          1 routes
         iddn:          3 routes
```

show route forwarding-table extensive

```
user@switch> show route forwarding-table summary

Routing table: default.inet [Index 0]
Internet:

Destination:  default
  Route type: user
  Route reference: 2                    Route interface-index: 0
  Flags: sent to PFE, rt nh decoupled
  Nexthop: 0:12:f2:21:cf:0
  Next-hop type: unicast                Index: 333      Reference: 5
  Next-hop interface: me0.0

Destination:  default
  Route type: permanent
  Route reference: 0                    Route interface-index: 0
```

```
  Flags: none
  Next-hop type: reject               Index: 36       Reference: 2

Destination:  0.0.0.0/32
  Route type: permanent
  Route reference: 0                  Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: discard              Index: 34       Reference: 1

Destination:  2.2.2.0/24
  Route type: interface
  Route reference: 0                  Route interface-index: 66
  Flags: sent to PFE
  Next-hop type: resolve              Index: 1309     Reference: 1
  Next-hop interface: ae0.0

Destination:  2.2.2.0/32
  Route type: destination
  Route reference: 0                  Route interface-index: 66
  Flags: sent to PFE
  Nexthop: 2.2.2.0
  Next-hop type: receive              Index: 1307     Reference: 1
  Next-hop interface: ae0.0

Destination:  2.2.2.1/32
  Route type: destination
  Route reference: 0                  Route interface-index: 66
  Flags: sent to PFE
  Nexthop: 0:21:59:cc:89:c0
  Next-hop type: unicast              Index: 1320     Reference: 1
  Next-hop interface: ae0.0

Destination:  2.2.2.2/32
  Route type: interface
  Route reference: 0                  Route interface-index: 0
  Flags: sent to PFE
  Nexthop: 2.2.2.2
  Next-hop type: local                Index: 1308     Reference: 2

Destination:  2.2.2.2/32
  Route type: destination
  Route reference: 0                  Route interface-index: 66
  Flags: none
  Nexthop: 2.2.2.2
  Next-hop type: local                Index: 1308     Reference: 2

Destination:  2.2.2.255/32
  Route type: destination
  Route reference: 0                  Route interface-index: 66
  Flags: sent to PFE
  Nexthop: 2.2.2.255
  Next-hop type: broadcast            Index: 1306     Reference: 1
  Next-hop interface: ae0.0
```

show route forwarding-table ccc

```
user@switch> show route forwarding-table ccc ge-0/0/0.10
Routing table: default.mpls
MPLS:
Destination        Type RtRef Next hop           Type Index NhRef Netif
ge-0/0/0.10  (CCC) user     0 3.3.3.2           Push 300112  1343     2 ae1.0
```

show route forwarding-table family

```
user@switch> show route forwarding-table family mpls

Routing table: default.mpls
MPLS:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    dscd    50     1
0                  user     0                    recv    49     3
1                  user     0                    recv    49     3
2                  user     0                    recv    49     3
299776             user     0                   Pop    1334     2 ge-0/0/0.10
299792             user     0                   Pop    1339     2 ge-0/0/0.14
299808             user     0                   Pop    1341     2 ge-0/0/0.2
299824             user     0                   Pop    1344     2 ge-0/0/0.11
299840             user     0                   Pop    1345     2 ge-0/0/0.13
299856             user     0                   Pop    1346     2 ge-0/0/0.18
299872             user     0                   Pop    1347     2 ge-0/0/0.16
299888             user     0                   Pop    1348     2 ge-0/0/0.7
299904             user     0                   Pop    1349     2 ge-0/0/0.20
299920             user     0                   Pop    1350     2 ge-0/0/0.19
299936             user     0                   Pop    1351     2 ge-0/0/0.17
299952             user     0                   Pop    1352     2 ge-0/0/0.9
299968             user     0                   Pop    1353     2 ge-0/0/0.1
299984             user     0                   Pop    1354     2 ge-0/0/0.12
300000             user     0                   Pop    1355     2 ge-0/0/0.8
300016             user     0                   Pop    1356     2 ge-0/0/0.4
300032             user     0                   Pop    1357     2 ge-0/0/0.5
300048             user     0                   Pop    1358     2 ge-0/0/0.3
300064             user     0                   Pop    1359     2 ge-0/0/0.15
ge-0/0/0.1  (CCC)  user     0 3.3.3.2           Push 300064  1340     2 ae1.0
ge-0/0/0.2  (CCC)  user     0 3.3.3.2           Push 299872  1328     2 ae1.0
ge-0/0/0.3  (CCC)  user     0 3.3.3.2           Push 299792  1323     2 ae1.0
ge-0/0/0.4  (CCC)  user     0 3.3.3.2           Push 300016  1337     2 ae1.0
ge-0/0/0.5  (CCC)  user     0 3.3.3.2           Push 299824  1325     2 ae1.0
ge-0/0/0.7  (CCC)  user     0 3.3.3.2           Push 299920  1331     2 ae1.0
ge-0/0/0.8  (CCC)  user     0 3.3.3.2           Push 299840  1326     2 ae1.0
ge-0/0/0.9  (CCC)  user     0 3.3.3.2           Push 299888  1329     2 ae1.0
ge-0/0/0.10  (CCC) user     0 3.3.3.2           Push 300112  1343     2 ae1.0
ge-0/0/0.11  (CCC) user     0 3.3.3.2           Push 299776  1322     2 ae1.0
ge-0/0/0.12  (CCC) user     0 3.3.3.2           Push 299952  1333     2 ae1.0
ge-0/0/0.13  (CCC) user     0 3.3.3.2           Push 300096  1342     2 ae1.0
ge-0/0/0.14  (CCC) user     0 3.3.3.2           Push 299984  1335     2 ae1.0
ge-0/0/0.15  (CCC) user     0 3.3.3.2           Push 299936  1332     2 ae1.0
ge-0/0/0.16  (CCC) user     0 3.3.3.2           Push 299808  1324     2 ae1.0
ge-0/0/0.17  (CCC) user     0 3.3.3.2           Push 300000  1336     2 ae1.0
ge-0/0/0.18  (CCC) user     0 3.3.3.2           Push 300032  1338     2 ae1.0
ge-0/0/0.19  (CCC) user     0 3.3.3.2           Push 299904  1330     2 ae1.0
ge-0/0/0.20  (CCC) user     0 3.3.3.2           Push 299856  1327     2 ae1.0
```

show route forwarding-table label

```
user@switch> show route forwarding-table label 29976

Routing table: default.mpls
MPLS:
Destination        Type RtRef Next hop           Type Index NhRef Netif
299776             user     0                   Pop    1334     2 ge-0/0/0.10
```

show route forwarding-table matching

```
user@switch> show route forwarding-table matching 3

Routing table: default.inet
Internet:
```

show route forwarding-table multicast

```
user@switch> show route forwarding-table multicast

Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
224.0.0.0/4        perm     1                    mdsc    35     1
224.0.0.1/32       perm     0 224.0.0.1          mcst    31     3
224.0.0.5/32       user     1 224.0.0.5          mcst    31     3

Routing table: __master.anon__.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
224.0.0.0/4        perm     0                    mdsc  1289     1
224.0.0.1/32       perm     0 224.0.0.1          mcst  1285     1

Routing table: default.inet6
Internet6:
Destination        Type RtRef Next hop           Type Index NhRef Netif
ff00::/8           perm     0                    mdsc    43     1
ff02::1/128        perm     0 ff02::1            mcst    39     1
```

show rsvp interface

Syntax
show rsvp interface
<brief | detail | extensive>
<link-management>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show rsvp interface
<brief | detail | extensive>
<link-management>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display the status of Resource Reservation Protocol (RSVP)-enabled interfaces and packet statistics.

Options
none—Display standard information about the status of RSVP-enabled interfaces and packet statistics.

brief | detail | extensive | link-management—(Optional) Display the specified level of output.

link-management—(Optional) Use the link-management option to display the control peers and corresponding TE-link information created by the Link Management Protocol (LMP).

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show rsvp interface brief on page 3209
show rsvp interface detail on page 3209
show rsvp interface extensive on page 3209
show rsvp interface link-management on page 3210

Output Fields
Table 428 on page 3206 lists the output fields for the **show rsvp interface** command. Output fields are listed in the approximate order in which they appear.

Table 428: show rsvp interface Output Fields

Field Name	Field Description	Level of Output
RSVP interface	Number of interfaces on which RSVP is active. Each interface has one line of output.	All levels
Interface	Name of the interface.	All levels
Index	Index of the interface.	**detail**

Table 428: show rsvp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
State	State of the interface. • **Disabled**—No traffic engineering information is displayed. • **Down**—Interface is not operational. • **Enabled**—Displays traffic engineering information. • **Up**—Interface is operational.	All levels
NoAuthentication	Interface does not support RSVP authentication.	**detail**
NoAggregate	Interface does not support refresh reduction.	**detail**
NoReliable	Interface does not support refresh reduction message ID extension.	**detail**
NoLinkProtection	Interface does not support link protection.	**detail**
HelloInterval	Frequency at which RSVP hellos are sent on this interface (in seconds).	**detail**
Address	IP address of the local interface.	**detail**
Active control channel	Next-hop link address to transmit messages.	None specified
TElink	Traffic-engineered links that are managed by the peer they are associated with.	None specified
Active resv	Number of reservations that are actively reserving bandwidth on the interface.	All levels
PreemptionCnt	Number of times an RSVP session was preempted on this interface.	**detail**
Update threshold	Percentage change in reserved bandwidth to trigger an IGP update.	**detail**
Subscription	User-configured subscription factor.	All levels
bc *number*	Bandwidth allocated for the specified bandwidth constraint.	**extensive**
ct *number*	Bandwidth allocated for the specified class type.	**extensive**
Static BW	Total interface bandwidth, in bps.	All levels
Available BW	Amount of bandwidth that RSVP is allowed to reserve, in bps. It is equal to (static bandwidth * subscription factor).	al levels
Reserved BW	Currently reserved bandwidth, in bps.	All levels
SoftPreemptionCnt	Number of times a soft preemption occurred on this interface. This number is not included in the **PreemptionCnt** value.	**detail**
Overbooked BW	Currently overbooked bandwidth, in bps, by class type (**ct0** through **ct3**).	**detail**

Table 428: show rsvp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
Highwater mark	Highest bandwidth that has ever been reserved on this interface, in bps.	**brief**
PacketType	Type of RSVP packet.	**detail**
Total Sent	Total number of packets sent.	**detail**
Total Received	Total number of packets received since RSVP was enabled.	**detail**
Last 5 seconds Sent	Number of packets sent in the last 5 seconds.	**detail**
Last 5 seconds Received	Number of packets received in the last 5 seconds.	**detail**
Path	Statistics about Path messages, which are sent from the RSVP sender along the data paths and store path state information in each node along the path.	**detail**
PathErr	Statistics about PathErr messages, which are advisory messages that are sent upstream to the sender.	**detail**
PathTear	Statistics about PathTear messages, which remove path states and dependent reservation states in any routers along a path.	**detail**
Resv	Statistics about Resv messages, which are sent from the RSVP receiver along the data paths and store reservation state information in each node along the path.	**detail**
ResvErr	Statistics about ResvErr messages, which are advisory messages that are sent when an attempt to establish a reservation fails.	**detail**
ResvTear	Statistics about ResvTear messages, which remove reservation states along a path.	**detail**
Hello	Number of RSVP hello packets that have been sent to and received from the neighbor.	**detail**
Ack	Acknowledge message for refresh reductions.	**detail**
Srefresh	Summary refresh messages.	**detail**
EndtoEnd RSVP	Statistics for the number of end-to-end RSVP messages sent.	**detail**
Queue	CoS transmit queue number and its associated forwarding class designation.	**extensive**
TxRate	Configured bandwidth in Mbps and configured bandwidth as a percentage of the specified queue.	**extensive**
Priority	Weight of the queue relative to other configured queues, in percentage.	**extensive**

Table 428: show rsvp interface Output Fields *(continued)*

Field Name	Field Description	Level of Output
queue-priority-value	**Low**, **High**, **None**, or **Exact**. **None** indicates no rate limiting. **Exact** indicates the queue transmits at the configured rate only.	**extensive**

show rsvp interface brief

```
user@host> show rsvp interface brief
RSVP interface: 1 active
                    Active Subscr- Static      Available   Reserved    Highwater
Interface   State resv   iption  BW          BW          BW          mark
de0.0     Up         1    23%    10Mbps      989.992kbps 1.31Mbps    1.31Mbps
```

show rsvp interface detail

```
user@host> show rsvp interface detail
so-0/1/1.0  Index 6, State: Ena/Up
  NoAuthentication, NoAggregate, NoReliable, NoLinkProtection
  HelloInterval 3(second)
  Address 192.168.207.29, 10.255.245.194
  ActiveResv 0, PreemptionCnt 0, Update threshold 10%
  Subscription 100%, StaticBW 155.52Mbps, AvailableBW 155.52Mbps
ReservedBW [0] 155Mbps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
SoftPreemptionCnt1
OverbookedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 155Mbps[5] 0bps[6] 0bps[7] 0bps
PacketType              Total                 Last 5 seconds
                    Sent      Received        Sent      Received
  Path                16             0           1             0
  PathErr              0             0           0             0
  PathTear             1             0           0             0
  Resv                 0            11           0             1
  ResvErr              0             0           0             0
  ResvTear             0             0           0             0
  Hello               66            67           1             1
  Ack                  0             0           0             0
  Srefresh             0             0           0             0
  EndtoEnd RSVP        0             0           0             0
...
```

show rsvp interface extensive

```
user@host> show rsvp interface extensive
so-1/0/0.0 Index 72, State Ena/Up
  NoAuthentication, NoAggregate, NoReliable, NoLinkProtection
  HelloInterval 9(second)
  Address 192.168.213.22, 10.255.240.175
  ActiveResv 1, PreemptionCnt 0, Update threshold 10%
  Subscription 100%,
  bc0 = (ct0+ct1+ct2+ct3), StaticBW 622.08Mbps
  bc1 = (ct1+ct2+ct3), StaticBW 466.56Mbps
  bc2 = (ct2+ct3), StaticBW 311.04Mbps
  bc3 = ct3, StaticBW 155.52Mbps
  ct0: StaticBW 155.52Mbps, AvailableBW 522.08Mbps
  ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
  ct1: StaticBW 155.52Mbps, AvailableBW 366.56Mbps
  ReservedBW [0] 100Mbps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps

  ct2: StaticBW 155.52Mbps, AvailableBW 311.04Mbps
  ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
  ct3: StaticBW 155.52Mbps, AvailableBW 155.52Mbps
  ReservedBW [0] 0bps[1] 0bps[2] 0bps[3] 0bps[4] 0bps[5] 0bps[6] 0bps[7] 0bps
  Queue          TxRate             Priority Exact
    0            155.52Mbps         25%      Low
```

```
    1            155.52Mbps          25%       Low
    2            155.52Mbps          25%       Low
    3            155.52Mbps          25%       Low
```

show rsvp interface link-management

```
user@host> show rsvp interface link-management
RSVP interface: 2 active
PEER-C  State: Up
Active Control Channel: so-0/1/0.0

  TElink: TElnk1, Link ID: 37811
  ActiveResv 0, PreemptionCnt 0
  StaticBW 155.52Mbps, ReservedBW: 0bps, AvailableBW: 155.52Mbps

  TElink: TElnk2, Link ID: 37808
  ActiveResv 1, PreemptionCnt 0
  StaticBW 155.52Mbps, ReservedBW: 0bps, AvailableBW: 155.52Mbps

PEER-B  State: Up
Active Control Channel: so-1/0/0.0

  TElink: TElnkAB1, Link ID: 1598
  ActiveResv 0, PreemptionCnt 0
  StaticBW 622.08Mbps, ReservedBW: 0bps, AvailableBW: 622.08Mbps

  TElink: TElnkAB2, Link ID: 1597
  ActiveResv 0, PreemptionCnt 0
  StaticBW 622.08Mbps, ReservedBW: 0bps, AvailableBW: 622.08Mbps
```

show rsvp neighbor

Syntax
show rsvp neighbor
<brief | detail>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show rsvp neighbor
<brief | detail>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display Resource Reservation Protocol (RSVP) neighbors that were discovered dynamically during the exchange of RSVP packets.

Options
none—Display standard information about RSVP neighbors.

brief | detail—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show rsvp neighbor on page 3215
show rsvp neighbor detail on page 3215

Output Fields
Table 429 on page 3211 lists the output fields for the **show rsvp neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 429: show rsvp neighbor Output Fields

Field Name	Field Description	Level of Output
RSVP neighbor	Number of neighbors that the routing device has learned of. Each neighbor has one line of output.	All levels
via	Name of the interface where the neighbor has been detected. In the case of generalized MPLS (GMPLS) LSPs, the name of the peer where the neighbor has been detected.	**detail**
Address	Address of a learned neighbor.	All levels
Idle	Length of time the neighbor has been idle, in seconds.	All levels
Up/Dn	Number of neighbor up or down transitions detected by RSVP hello packets. If the up count is 1 greater than the down count, the neighbor is currently up. Otherwise, the neighbor is down. Neighbors that do not support RSVP hello packets are not reported as up or down.	All levels

Table 429: show rsvp neighbor Output Fields *(continued)*

Field Name	Field Description	Level of Output
Up cnt and Down cnt	Number of neighbor up or down transitions detected by RSVP hello packets. If the up count is 1 greater than the down count, the neighbor is currently up. Otherwise, the neighbor is down. Neighbors that do not support RSVP hello packets are not reported as up or down.	**detail**

status	State of the RSVP neighbor:
	• **Up**—Routing device can detect RSVP Hello messages from the neighbor.
	• **Down**—Routing device has received one of the following indications:
	  ◦ Communication failure from the neighbor.
	  ◦ Communication from IGP that the neighbor is unavailable.
	  ◦ Change in the sequence numbers in the RSVP Hello messages sent by the neighbor.
	• **Restarting**—RSVP neighbor is unavailable and might be restarting. The neighbor remains in this state until it has restarted or is declared dead. This state is possible only when graceful restart is enabled.
	• **Restarted**—RSVP neighbor has restarted and is undergoing state recovery (graceful restart) procedures.
	• **Dead**—Routing device has lost all communication with the RSVP neighbor. Any RSVP sessions with that neighbor are torn down.
		detail

LastChange	Time elapsed since the neighbor state changed either from up to down or from down to up. The format is ***hh*:*mm*:*ss***.	All levels
Last changed time	Time elapsed since the neighbor state changed either from up to down or from down to up.	**detail**
HelloInt	Frequency at which RSVP hellos are sent on this interface (in seconds).	All levels
HelloTx/Rx	Number of hello packets sent to and received from the neighbor.	All levels
Hello	Number of RSVP hello packets that have been sent to and received from the neighbor.	**detail**
Message received	Number of Path and Resv messages that this routing device has received from the neighbor.	**detail**
Remote Instance	Identification provided by the remote routing device during Hello message exchange.	**detail**
Local Instance	Identification sent to the remote routing device during Hello message exchange.	**detail**

Table 429: show rsvp neighbor Output Fields *(continued)*

Field Name	Field Description	Level of Output
Refresh reduction	Measure of processing overhead requests of refresh messages. Refresh reduction extensions improve routing device performance by reducing the process overhead, thus increasing the number of LSPs a routing device can support. **Refresh reduction** can have the following values: • **operational**—All four RSVP refresh reduction extensions—message ack, bundling, summary refresh, and staged refresh timer—are functional between the two neighboring routing devices. For a detailed explanation of these extensions, see RFC 2961. • **incomplete**—Some RSVP refresh reduction extensions are functional between the two neighboring routing devices. • **no operational**—Either the refresh reduction feature has been turned off, or the remote routing device cannot support the refresh reduction extensions.	**detail**
Remote end	Neighboring routing device's status with regard to refresh reduction: • **enabled**—Remote routing device has requested refresh reduction during RSVP message exchanges. • **disabled**—Remote routing device does not require refresh reduction.	**detail**
Ack-extension	An RSVP refresh reduction extension: • **enabled**—Both local and remote routing devices support the ack-extension (RFC 2961). • **disabled**—Remote routing device does not support the ack-extension.	**detail**
Link protection	Status of the MPLS fast reroute mechanism that protects traffic from link failure: • **enabled**—Link protection feature has been turned on, protecting the neighbor with a bypass LSP. • **disabled**—No link protection feature has been enabled for this neighbor.	**detail**
LSP name	Name of the bypass LSP.	**detail**
Bypass LSP	Status of the bypass LSP. It can have the following values: • **does not exist**—Bypass LSP is not available. • **connecting**—Routing device is in the process of establishing a bypass LSP, and the LSP is not available for link protection at the moment. • **operational**—Bypass LSP is up and running. • **down**—Bypass LSP has gone down, with the most probable cause a node or a link failure on the bypass path.	**detail**
Backup routes	Number of user LSPs (or routes) that are being protected by a bypass LSP (before link failure).	**detail**
Backup LSPs	Number of LSPs that have been temporarily established to maintain traffic by refreshing the downstream LSPs during link failure (not a one-to-one correspondence).	**detail**
Bypass explicit route	Explicit route object's (ERO) path that is taken by the bypass LSP.	**detail**

Table 429: show rsvp neighbor Output Fields *(continued)*

Field Name	Field Description	Level of Output
Restart time	Length of time a neighbor waits to receive a Hello from the restarting node before declaring the node dead and deleting the states (in milliseconds).	**detail**
Recovery time	Length of time during which the restarting node attempts to recover its lost states with help from its neighbors (in milliseconds). Recovery time is advertised by the restarting node to its neighbors, and applies to nodal faults. The restarting node considers its graceful restart complete after this time has elapsed.	**detail**

show rsvp neighbor

```
user@host> show rsvp neighbor
RSVP neighbor: 2 learned
Address            Idle Up/Dn LastChange HelloInt HelloTx/Rx
192.168.207.203       0  3/2       13:01        3   366/349
192.168.207.207       0  1/0       22:49        3   448/448
```

show rsvp neighbor detail

```
user@host> show rsvp neighbor detail
RSVP neighbor: 2 learned
Address: 192.168.207.203   via: ecstasyl status: Up
   Last changed time: 28:47, Idle: 0 sec, Up cnt: 3, Down cnt: 2
   Message received: 632
   Hello: sent 673, received 656, interval 3 sec
   Remote instance: 0x6432838a, Local instance: 0x74b72e36
   Refresh reduction:  operational
     Remote end: enabled,  Ack-extension: enabled
   Link protection:  enabled
     LSP name: Bypass_to_192.168.207.203
     Bypass LSP: operational,  Backup routes: 1,  Backup LSPs: 0
     Bypass explicit route: 192.168.207.207 192.168.207.224
  Restart time: 60000 msec, Recovery time: 0 msec
```

show rsvp session

Syntax

show rsvp session
<brief | detail | extensive | terse>
<bidirectional | unidirectional>
<down | up>
<interface *interface-name*>
<*lsp-type*>
<name *session-name*>
<*session-type*>
<statistics>
<te-link *te-link*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display information about Resource Reservation Protocol (RSVP) sessions.

Options

none—Display standard information about all RSVP sessions.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

bidirectional | unidirectional—(Optional) Display information about bidirectional or unidirectional RSVP sessions only, respectively.

down | up—(Optional) Display only LSPs that are inactive or active, respectively.

interface *interface-name*—(Optional) Display RSVP sessions for the specified interface only.

lsp-type —(Optional) Display information about RSVP sessions with regard to LSPs:

- **bypass**—Sessions used for bypass LSPs.
- **lsp**—Sessions used to set up LSPs.
- **nolsp**—Sessions not used to set up LSPs.

name *session-name*—(Optional) Display information about the named session.

session-type—(Optional) Display information about a particular session type:

- **egress**—Sessions that terminate on this switch.
- **ingress**—Sessions that originate from this switch.
- **transit**—Sessions that transit through this switch.

statistics—(Optional) Display packet statistics.

te-link *te-link*—(Optional) Display sessions with reservations on the specified traffic-engineered link name.

Required Privilege Level

view

Related Documentation

- Example: Configuring MPLS on J-EX Series Switches on page 3071
- Configuring MPLS on Provider Edge Switches (CLI Procedure)
- Configuring MPLS on Provider Switches (CLI Procedure) on page 3102
- *Junos OS MPLS Applications Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

List of Sample Output

show rsvp session on page 3218
show rsvp session statistics on page 3219
show rsvp session detail on page 3219
show rsvp session extensive on page 3219

Output Fields

Table 430 on page 3217 describes the output fields for the **show rsvp session** command. Output fields are listed in the approximate order in which they appear.

Table 430: show rsvp session Output Fields

Field Name	Field Description	Level of Output
Ingress RSVP	Information about ingress RSVP sessions.	**detail**
Ingress RSVP	Information about ingress RSVP sessions. Each session has one line of output.	All levels
Egress RSVP	Information about egress RSVP sessions.	All levels
Transit RSVP	Information about the transit RSVP sessions.	All levels
To	Destination (egress switch) of the session.	All levels
From	Source (ingress switch) of the session.	All levels
State	State of the path: **Up**, **Down**, or **AdminDn**. **AdminDn** indicates that the LSP is being taken down gracefully.	All levels
Address	Destination (egress switch) of the LSP.	**detail**
LSPstate	State of the LSP that is being handled by this RSVP session. It can be either **Up**, **Dn** (down), or **AdminDn**. **AdminDn** indicates that the LSP is being taken down gracefully.	**brief, detail**
Rt	Number of active routes (prefixes) that have been installed in the routing table. For ingress RSVP sessions, the routing table is the primary IPv4 table (**inet.0**). For transit and egress RSVP sessions, the routing table is the primary MPLS table **mpls.0**).	**brief**
ActiveRoute	Number of active routes (prefixes) that have been installed in the forwarding table. For ingress RSVP sessions, the forwarding table is the primary IPv4 table (**inet.0**). For transit and egress RSVP sessions, the forwarding table is the primary MPLS table (**mpls.0**).	**detail**
LSPname	Name of the LSP.	**brief, detail**

Table 430: show rsvp session Output Fields *(continued)*

Field Name	Field Description	Level of Output
LSPpath	Indicates whether the RSVP session is for the primary or secondary LSP path. **LSPpath** can be either **primary** or **secondary** and can be displayed on the ingress, egress, and transit switches. **LSPpath** can also indicate when a graceful LSP deletion has been triggered.	**detail**
Recovery label received	(When LSP is bidirectional) Label the upstream node suggests for use in the Resv message that is sent.	**detail**
Recovery label sent	(When LSP is bidirectional) Label the downstream node suggests for use in its Resv messages that is returned.	**detail**
Suggested label received	(When LSP is bidirectional) Label the upstream node suggests for use in the Resv message that is sent.	**detail**
Suggested label sent	(When LSP is bidirectional) Label the downstream node suggests for use in its Resv message that is returned.	**detail**
Resv style or **Style**	RSVP reservation style. This field consists of two parts. The first is the number of active reservations. The second is the reservation style, which can be **FF** (fixed filter), **SE** (shared explicit), or **WF** (wildcard filter).	**brief detail**
Label in	Incoming label for this LSP.	**brief, detail**
Label out	Outgoing label for this LSP.	**brief, detail**
Time left	Number of seconds remaining in the lifetime of the reservation.	**brief, detail**
Since	Date and time when the RSVP session was initiated.	**detail**
Tspec	Sender's traffic specification, which describes the sender's traffic parameters.	**detail**
Port number	Protocol ID and sender/receiver port used in this RSVP session.	**detail**
Creating backup LSP, link down	A **link down** event occurred, and traffic is being switched over to the bypass LSP.	**extensive**
Deleting backup LSP, protected LSP restored	Link has come back up and the LSP has been restored. Because the backup LSP is no longer needed, it is deleted.	**extensive**
PATH rcvfrom	Address of the previous-hop (upstream) switch or client, interface the neighbor used to reach this switch, and number of packets received from the upstream neighbor.	**detail**

show rsvp session

```
user@switch> show rsvp session
Ingress RSVP: 1 sessions
To              From            State   Rt Style Labelin Labelout LSPname
10.255.245.214  10.255.245.212  AdminDn  0  1 FF       -    22293 LSP Bidir
Total 1 displayed, Up 1, Down 0
```

```
Egress RSVP: 2 sessions
To              From            State Rt Style Labelin Labelout LSPname
10.255.245.194  10.255.245.195  Up     0  1 FF   39811        - Gpro3-ba Bidir
10.255.245.194  10.255.245.195  Up     0  1 FF       3        - pro3-ba
Total 2 displayed, Up 2, Down 0

Transit RSVP: 1 sessions
To              From            State Rt Style Labelin Labelout LSPname
10.255.245.198  10.255.245.197  Up     0  1 SE  100000        3 pro3-de
Total 1 displayed, Up 1, Down 0
```

show rsvp session statistics

```
user@switch> show rsvp session statistics
Ingress RSVP: 2 sessions
To              From           State   Packets      Bytes    LSPname
10.255.245.24   10.255.245.22     Up         0          0      pro3-bd
10.255.245.24   10.255.245.22     Up     44868    2333136  pro3-bd-2
Total 2 displayed, Up 2, Down 0
Egress RSVP: 2 sessions
To              From           State   Packets      Bytes    LSPname
10.255.245.22   10.255.245.24     Up         0          0    pro3-db
10.255.245.22   10.255.245.24     Up         0          0  pro3-db-2
Total 2 displayed, Up 2, Down 0
Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

show rsvp session detail

```
user@switch> show rsvp session detail
Ingress RSVP: 1 sessions
1.1.1.1
  From: 2.2.2.2, LSPstate: Up, ActiveRoute: 0
  LSPname: to-a, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: -, Label out: 3
  Time left:    -, Since: Fri Mar 26 18:42:42 2004
  Tspec: rate 300kbps size 300kbps peak Infbps m 20 M 1500
  DiffServ info: diffServ-TE LSP, bandwidth: <ct1 300kbps>
  Port number: sender 1 receiver 15876 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  PATH sentto: 192.168.37.16 (t1-0/2/1.0) 1 pkt
```

show rsvp session extensive

```
user@switch> show rsvp session extensive

8.8.8.8
  From: 9.9.9.9, LSPstate: Up, ActiveRoute: 0
  LSPname: lsp_to_240, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 322832
  Resv style: 1 FF, Label in: -, Label out: 322832
  Time left:    -, Since: Thu Feb 26 16:25:39 2009
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 1500
  Port number: sender 2 receiver 44542 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  Path MTU: received 1500
  PATH sentto: 3.3.3.2 (xe-0/1/0.0) 238 pkts
  RESV rcvfrom: 3.3.3.2 (xe-0/1/0.0) 234 pkts
```

```
Explct route: 3.3.3.2 4.4.4.2
```

show rsvp session

Syntax

show rsvp session
<brief | detail | extensive | terse>
<bidirectional | unidirectional>
<bypass>
<down | up>
<interface *interface-name*>
<logical-system (all | *logical-system-name*)>
<*lsp-type*>
<name *session-name*>
<p2mp>
<*session-type*>
<statistics>
<te-link *te-link*>

Syntax (J-EX Series Switch)

show rsvp session
<brief | detail | extensive | terse>
<bidirectional | unidirectional>
<bypass>
<down | up>
<interface *interface-name*>
<*lsp-type*>
<name *session-name*>
<p2mp>
<*session-type*>
<statistics>
<te-link *te-link*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display information about Resource Reservation Protocol (RSVP) sessions.

Options

none—Display standard information about all RSVP sessions.

brief | detail | extensive | terse—(Optional) Display the specified level of output.

bidirectional | unidirectional—(Optional) Display information about bidirectional or unidirectional RSVP sessions only, respectively.

bypass—(Optional) Display RSVP sessions for bypass LSPs.

down | up—(Optional) Display only LSPs that are inactive or active, respectively.

interface *interface-name*—(Optional) Display RSVP sessions for the specified interface only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

lsp-type—(Optional) Display information about RSVP sessions with regard to LSPs:

- **bypass**—Sessions used for bypass LSPs.

- **lsp**—Sessions used to set up LSPs.
- **nolsp**—Sessions not used to set up LSPs.

name *session-name*—(Optional) Display information about the named session.

p2mp—(Optional) Display point-to-multipoint information.

session-type—(Optional) Display information about a particular session type:

- **egress**—Sessions that terminate on this routing device.
- **ingress**—Sessions that originate from this routing device.
- **transit**—Sessions that transit through this routing device.

statistics—(Optional) Display packet statistics.

te-link *te-link*—(Optional) Display sessions with reservations on the specified TE link.

Required Privilege Level view

Related Documentation
- **clear rsvp session on page 3142**

List of Sample Output
show rsvp session on page 3226
show rsvp session statistics on page 3226
show rsvp session detail on page 3226
show rsvp session detail (Path MTU Output Field) on page 3227
show rsvp session detail (GMPLS) on page 3227
show rsvp session extensive on page 3227
show rsvp session p2mp on page 3228

Output Fields Table 431 on page 3222 describes the output fields for the **show rsvp session** command. Output fields are listed in the approximate order in which they appear.

Table 431: show rsvp session Output Fields

Field Name	Field Description	Level of Output
Ingress RSVP	Information about ingress RSVP sessions.	**detail**
Ingress RSVP	Information about ingress RSVP sessions. Each session has one line of output.	All levels
Egress RSVP	Information about egress RSVP sessions.	All levels
Transit RSVP	Information about the transit RSVP sessions.	All levels
P2MP name	(Appears only when the **p2mp** option is specified). Name of the point-to-multipoint LSP path.	All levels
P2MP branch count	(Appears only when the **p2mp** option is specified). Number of LSPs receiving packets from the point-to-multipoint LSP.	All levels

Table 431: show rsvp session Output Fields *(continued)*

Field Name	Field Description	Level of Output
To	Destination (egress routing device) of the session.	All levels
From	Source (ingress routing device) of the session.	All levels
State	State of the path: **Up**, **Down**, or **AdminDn**. **AdminDn** indicates that the LSP is being taken down gracefully.	All levels
Address	Destination (egress routing device) of the LSP.	**detail**
From	Source (ingress routing device) of the session.	**detail**
LSPstate	State of the LSP that is being handled by this RSVP session. It can be either **Up**, **Dn** (down), or **AdminDn**. **AdminDn** indicates that the LSP is being taken down gracefully.	**brief detail**
Rt	Number of active routes (prefixes) that have been installed in the routing table. For ingress RSVP sessions, the routing table is the primary IPv4 table (**inet.0**). For transit and egress RSVP sessions, the routing table is the primary MPLS table **mpls.0**).	**brief**
Active Route	Number of active routes (prefixes) that have been installed in the forwarding table. For ingress RSVP sessions, the forwarding table is the primary IPv4 table (**inet.0**). For transit and egress RSVP sessions, the forwarding table is the primary MPLS table (**mpls.0**).	**detail**
LSPname	Name of the LSP.	**brief detail**
LSPpath	Indicates whether the RSVP session is for the primary or secondary LSP path. **LSPpath** can be either **primary** or **secondary** and can be displayed on the ingress, egress, and transit routing devices. **LSPpath** can also indicate when a graceful LSP deletion has been triggered.	**detail**
Bypass	(Egress routing device) Destination address for the bypass LSP.	**detail**
Bidir	(When LSP is bidirectional) LSP will allow data to travel in both directions between GMPLS devices.	**detail**
Bidirectional	(When LSP is bidirectional) LSP will allow data to travel both ways between GMPLS devices.	**detail**
Upstream label in	(When LSP is bidirectional) Incoming label for reverse direction traffic for this LSP.	**detail**
Upstream label out	(When LSP is bidirectional) Outgoing label for reverse direction traffic for this LSP.	**detail**
Recovery label received	(When LSP is bidirectional) Label the upstream node suggests for use in the Resv message that is sent.	**detail**

Table 431: show rsvp session Output Fields *(continued)*

Field Name	Field Description	Level of Output
Recovery label sent	(When LSP is bidirectional) Label the downstream node suggests for use in its Resv messages that is returned.	**detail**
Suggested label received	(When LSP is bidirectional) Label the upstream node suggests for use in the Resv message that is sent.	**detail**
Suggested label sent	(When LSP is bidirectional) Label the downstream node suggests for use in its Resv message that is returned.	**detail**
Resv style or **Style**	RSVP reservation style. This field consists of two parts. The first is the number of active reservations. The second is the reservation style, which can be **FF** (fixed filter), **SE** (shared explicit), or **WF** (wildcard filter).	**brief detail**
Label in	Incoming label for this LSP.	**brief detail**
Label out	Outgoing label for this LSP.	**brief detail**
Time left	Number of seconds remaining in the lifetime of the reservation.	**brief detail**
Since	Date and time when the RSVP session was initiated.	**detail**
Tspec	Sender's traffic specification, which describes the sender's traffic parameters.	**detail**
DiffServ info	Indicates whether the LSP is a multiclass LSP (**multiclass diffServ-TE LSP**) or a Differentiated-Services-aware traffic engineering LSP (**diffServ-TE LSP**).	**detail**
bandwidth	Bandwidth for each class type (**ct0**, **ct1**, **ct2**, or **ct3**).	**detail**
Port number	Protocol ID and sender/receiver port used in this RSVP session.	**detail**
FastReroute desired	Fast reroute has been requested by the ingress routing device.	**detail**
Soft preemption desired	Soft preemption has been requested by the ingress routing device.	**detail**
FastReroute desired	(Data [not a bypass or backup] LSP when the protection scheme has been requested) Fast reroute (one-to-one backup) has been requested by the ingress routing device.	**detail extensive**
Link protection desired	(Data [not a bypass or backup] LSP when the protection scheme has been requested) Link protection (many-to-one backup) has been requested by the ingress routing device.	**detail extensive**
Node/Link protection desired	(Data [not a bypass or backup] LSP when the protection scheme has been requested) Node and link protection (many-to-one backup) has been requested by the ingress routing device.	**detail extensive**

Table 431: show rsvp session Output Fields *(continued)*

Field Name	Field Description	Level of Output
Type	LSP type: • **Link protected LSP**—LSP has been protected by link protection at the outgoing interface. The name of the bypass used is also listed here (**extensive**). • **Node/Link protected LSP**—LSP has been protected by node and link protection at the outgoing interface. The name of the bypass used is also listed here (**extensive**). • **Protection down**—LSP is not currently protected. • **Bypass LSP**—LSP that is used to protected one or more user LSPs in case of link failure. • **Backup LSP at Point-of-Local-Repair (PLR)**—LSP that has been temporarily established to protected a user LSP at the ingress of a failed link. • **Backup LSP at Merge Point (MP)**—LSP that has been temporarily established to protected a user LSP at the egress of a failed link.	**detail extensive**
New bypass	New bypass (the bypass name is also displayed) has been activated to protect the LSP.	**extensive**
Link protection up, using *bypass-name*	Link protection (the bypass name is also displayed) has been activated for the LSP.	**extensive**
Creating backup LSP, link down	A **link down** event occurred, and traffic is being switched over to the bypass LSP.	**extensive**
Deleting backup LSP, protected LSP restored	Link has come back up and the LSP has been restored. Because the backup LSP is no longer needed, it is deleted.	**extensive**
Path mtu	Displays the value of the path MTU received from the network (through signaling) and the value used for forwarding. This value is only displayed on ingress routing devices with the **allow-fragmentation** statement configured at the **[edit protocols mpls path-mtu]** hierarchy level. If there is a detour LSP, the path MTU for the detour is also displayed.	**detail**
PATH rcvfrom	Address of the previous-hop (upstream) routing device or client, interface the neighbor used to reach this routing device, and number of packets received from the upstream neighbor.	**detail**
Adspec	MTU signaled from the ingress routing device to the egress routing device by means of the adspec object.	**detail**
PATH sentto	Address of the next-hop (downstream) routing device or client, interface used to reach this neighbor (or peer-name in the GMPLS LSP case), and number of packets sent to the downstream routing device.	**detail**
Explct route	Explicit route for the session. Normally this value will be the same as that of record route. Differences indicate that path rerouting has occurred, typically during fast reroute.	**detail**

Table 431: show rsvp session Output Fields *(continued)*

Field Name	Field Description	Level of Output
Record route	Recorded route for the session, taken from the record route object. Normally this value will be the same as that of explct route. Differences indicate that path rerouting has occurred, typically during fast reroute.	**detail**

show rsvp session

```
user@host> show rsvp session
Ingress RSVP: 1 sessions
To              From            State   Rt Style Labelin Labelout LSPname
10.255.245.214  10.255.245.212  AdminDn  0  1 FF       -    22293 LSP Bidir
Total 1 displayed, Up 1, Down 0

Egress RSVP: 2 sessions
To              From            State Rt Style Labelin Labelout LSPname
10.255.245.194  10.255.245.195  Up     0  1 FF   39811        - Gpro3-ba Bidir
10.255.245.194  10.255.245.195  Up     0  1 FF       3        - pro3-ba
Total 2 displayed, Up 2, Down 0

Transit RSVP: 1 sessions
To              From            State Rt Style Labelin Labelout LSPname
10.255.245.198  10.255.245.197  Up     0  1 SE  100000        3 pro3-de
Total 1 displayed, Up 1, Down 0
```

show rsvp session statistics

```
user@host> show rsvp session statistics
Ingress RSVP: 2 sessions
To              From           State   Packets      Bytes    LSPname
10.255.245.24   10.255.245.22     Up         0          0      pro3-bd
10.255.245.24   10.255.245.22     Up     44868    2333136  pro3-bd-2
Total 2 displayed, Up 2, Down 0
Egress RSVP: 2 sessions
To              From           State   Packets      Bytes    LSPname
10.255.245.22   10.255.245.24     Up         0          0    pro3-db
10.255.245.22   10.255.245.24     Up         0          0  pro3-db-2
Total 2 displayed, Up 2, Down 0
Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

show rsvp session detail

```
user@host> show rsvp session detail
Ingress RSVP: 1 sessions
1.1.1.1
  From: 2.2.2.2, LSPstate: Up, ActiveRoute: 0
  LSPname: to-a, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 3
  Resv style: 1 FF, Label in: -, Label out: 3
  Time left:    -, Since: Fri Mar 26 18:42:42 2004
  Tspec: rate 300kbps size 300kbps peak Infbps m 20 M 1500
  DiffServ info: diffServ-TE LSP, bandwidth: <ct1 300kbps>
  Port number: sender 1 receiver 15876 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  PATH sentto: 192.168.37.16 (t1-0/2/1.0) 1 pkt
```

show rsvp session detail (Path MTU Output Field)

```
user@host> show rsvp session detail
Ingress RSVP: 1 sessions
10.255.245.3
  From: 10.255.245.5, LSPstate: Up, ActiveRoute: 3
  LSPname: to-c, LSPpath: Primary
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 100432
  Resv style: 1 FF, Label in: -, Label out: 100432
  Time left:    -, Since: Mon Aug 16 17:54:40 2006
  Tspec: rate 0bps size 0bps peak Infbps m 20 M 9192
  Port number: sender 1 receiver 57843 protocol 0
  FastReroute desired
  PATH rcvfrom: localclient
  Adspec: sent MTU 4470
  Path mtu: received 4470, using 4458 for forwarding
  PATH sentto: 192.168.37.89 (so-0/2/3.0) 11 pkts
  RESV rcvfrom: 192.168.37.89 (so-0/2/3.0) 10 pkts
  Explct route: 192.168.37.89
  Record route: <self> 192.168.37.89 192.168.37.87
    Detour is Up
    Detour Tspec: rate 0bps size 0bps peak Infbps m 20 M 9192
    Detour adspec: sent MTU 1512
    Path mtu: received 1512, using 1500 for forwarding
```

show rsvp session detail (GMPLS)

```
user@host> show rsvp session detail
Ingress RSVP: 1 sessions
192.168.4.1
  From: 192.168.1.1, LSPstate: Dn, ActiveRoute: 0
  LSPname: gmpls-r1-to-r3, LSPpath: Primary
  Bidirectional, Upstream label in: 21253, Upstream label out: -
  Suggested label received: -, Suggested label sent: 21253
  Recovery label received: -, Recovery label sent: -
  Resv style: 0 -, Label in: -, Label out: -
  Time left:    -, Since: Mon Aug 16 17:54:40 2006
  Tspec: rate 0bps size 0bps peak 155.52Mbps m 20 M 1500
  Port number: sender 2 receiver 46115 protocol 0
  PATH rcvfrom: localclient
  Adspec: sent MTU 1500
  PATH MTU: received 0
  PATH sentto: 10.35.1.5 (so-0/2/3.0) 11 pkts
  Explct route: 100.100.100.100 93.93.93.93
  Record route: <self> 100.100.100.100 93.93.93.93
  Total 1 displayed, Up 0, Down 1
  Egress RSVP: 0 sessions
  Total 0 displayed, Up 0, Down 0
  Transit RSVP: 0 sessions
  Total 0 displayed, Up 0, Down 0
```

show rsvp session extensive

```
user@host> show rsvp session extensive
10.255.245.13
    From: 10.255.245.48, LSPstate: Up, ActiveRoute: 0
    ....
    Link protection desired
    Type: Link protected LSP, using p2
       11 Feb  6 15:24:16 Backup LSP: Call was cleared by RSVP
       10 Feb  6 15:24:16 Backup LSP: Session preempted
        9 Feb  6 15:24:16 Deleting backup LSP, protected LSP restored
        8 Feb  6 15:23:22 Backup LSP: Up  192.168.208.117(Label=3)
        7 Feb  6 15:23:22 Backup LSP: Record Route:  192.168.208.117(Label=3)
        6 Feb  6 15:23:19 Backup LSP: Explicit Route: wrong delivery
```

```
        5 Feb  6 15:23:19 Creating backup LSP, link down
        4 Feb  6 12:36:03 Link protection up, using p2
        3 Feb  6 12:35:56 New bypass p2
        2 Feb  6 12:35:47 Bypass state down, p1[2 times]
        1 Feb  6 12:35:39 New bypass p1
```

show rsvp session p2mp

```
user@host> show rsvp session p2mp
Ingress RSVP: 3 sessions
P2MP name: p2mp-lsp1, P2MP branch count: 1
To              From            State Rt Style Labelin Labelout LSPname
10.255.245.34   10.255.245.25 Up     0 1 FF       -      100128 p2mp-branch-1
P2MP name: p2mp-lsp2, P2MP branch count: 1
To              From            State Rt Style Labelin Labelout LSPname
10.255.245.34   10.255.245.25 Up     0 1 FF       -           3 p2mp-st-br1
P2MP name: lsp-a_b, P2MP branch count: 1
Total 2 displayed, Up 2, Down 0

Egress RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit RSVP: 0 sessions
Total 0 displayed, Up 0, Down 0
```

show rsvp statistics

Syntax
show rsvp statistics
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show rsvp statistics

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display Resource Reservation Protocol (RSVP) packet and error statistics.

Options
none—Display RSVP packet and error statistics.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- **clear rsvp statistics on page 3144**

List of Sample Output
show rsvp statistics on page 3231

Output Fields
Table 432 on page 3229 describes the output fields for the **show rsvp statistics** command. Output fields are listed in the approximate order in which they appear.

Table 432: show rsvp statistics Output Fields

Field Name	Field Description
Packet Type	Statistics about different RSVP messages.
Total Sent	Total number of packets sent since RSVP was enabled.
Total Received	Total number of packets received since RSVP was enabled.
Last 5 seconds Sent	Total number of packets sent in the last 5 seconds.
Last 5 seconds Received	Number of packets received in the last 5 seconds.
Path	Statistics about Path messages, which are sent from the RSVP sender along the data paths and which store path state information in each node along the path.
PathErr	Statistics about PathErr messages, which are advisory messages that are sent upstream to the sender.
PathTear	Statistics about PathTear messages, which remove path states and dependent reservation states in any routing devices along a path.

Table 432: show rsvp statistics Output Fields *(continued)*

Field Name	Field Description
Resv FF	Statistics about fixed-filter reservation style messages, which consist of distinct reservations among explicit senders.
Resv WF	Statistics about wildcard-filter reservation style messages, which consist of shared reservations among wildcard senders.
Res SE	Statistics about shared-explicit reservation style messages, which consist of shared reservations among explicit senders.
ResvErr	Statistics about ResvErr messages, which are advisory messages that are sent when an attempt to establish a reservation fails.
ResvTear	Statistics about ResvTear messages, which remove reservation states along a path.
ResvConf	Statistics about ResvConfirm messages, which are responses to confirm a reservation request.
Ack	Acknowledge message for refresh reductions.
SRefresh	Summary refresh messages.
Hello	Number of RSVP hello packets that have been sent to and received from the neighbor.
EndtoEnd RSVP	Statistics for the number of End-to-end RSVP messages.
Errors	Statistics about errored RSVP packets.
Rcv pkt bad length	The packet was not processed because its length is inappropriate.
Rcv pkt unknown type	The packet is not one of the well-known RSVP types, as defined in RFC 2205, *Resource ReSerVation Protocol (RSVP)*.
Rcv pkt bad version	The packet is not an RSVP version 1 packet.
Rcv pkt auth fail	The packet failed authentication checks.
Rcv pkt bad checksum	The RSVP checksum check failed.
Rcv pkt bad format	General packet processing failed because the packet was badly formed.
Memory allocation fail	An internal resource failure occurred.
No path information	A reservation was received, but no sender is active.
Resv style conflict	The same session contains inconsistent reservation styles.
Port conflict	There were inconsistent port numbers for the same session.
Resv no interface	An interface for the receive reservation packets cannot be located.

Table 432: show rsvp statistics Output Fields *(continued)*

Field Name	Field Description
PathErr to client	Number of PathErr packets delivered to the local client.
ResvErr to client	Number of ResvErr packets delivered to the local client.
Path timeout	Number of times the sender timed out because the path was removed.
Resv timeout	Number of times the receiver timed out because the reservation was removed.
Message out-of-order	Records the number of RSVP incoming messages that are considered out of order. This is detected from the message ID object's sequence number.
Unknown ack msg	A neighboring routing device replies with an ACK object that contains an unknown message ID. This can indicate a message ID handshake problem.
Recv nack	A neighboring routing device explicitly rejects a message ID in a summary refresh message. This can happen if that neighbor has been rebooted. In this case, the routing device sends a regular RSVP refresh message to recover the state, and starts the message ID handshake process again.
Recv duplicated msg-id	Number of times the same message ID is used by two different RSVP messages. This duplication is usually caused when a neighboring routing device restarts.
No TE-link to recv Hop	Counter of packets discarded because a TE link was not found.
Rcv pkt disabled interface	Number of RSVP packets received on an interface that is not enabled for RSVP.
Transmit buffer full	Number of times the buffer for assembling an outgoing RSVP message was not large enough.
Transmit failure	Number of times the RSVP task failed to send out a packet.
Receive failure	Number of times the RSVP task failed to read an incoming packet.
P2MP RESV discarded by appl	Number of Resv messages discarded because the MPLS label is not valid for the P2MP LSP application.
Rate limit	Number of RSVP packets dropped due to rate limiting.
Err msg loop detected	Number of RSVP error messages that have looped back to their originator. This is detected by checking the error node address in the ERROR_SPEC object.

show rsvp statistics

```
user@host> show rsvp statistics
   PacketType              Total                 Last 5 seconds
                     Sent      Received        Sent      Received
   Path               355           408           0             0
   PathErr              2            13           0             0
   PathTear           101           139           0             0
   Resv FF              0             0           0             0
   Resv WF              0             0           0             0
   Resv SE            419           225           0             0
```

```
ResvErr              0             0           0             0
ResvTear             0            13           0             0
ResvConf             0             0           0             0
Ack                682          1414           0             0
SRefresh        395198        236030           5             2
Hello           578809        578221           4             4
EndtoEnd RSVP        0             0           0             0

Errors                          Total           Last 5 seconds
Rcv pkt bad length                  0                        0
Rcv pkt unknown type                0                        0
Rcv pkt bad version                 0                        0
Rcv pkt auth fail                   0                        0
Rcv pkt bad checksum                0                        0
Rcv pkt bad format                  0                        0
Memory allocation fail              0                        0
No path information                10                        0
Resv style conflict                 0                        0
Port conflict                       0                        0
Resv no interface                   0                        0
PathErr to client                  38                        0
ResvErr to client                   0                        0
Path timeout                        8                        0
Resv timeout                       57                        0
Message out-of-order                0                        0
Unknown ack msg                  2978                        0
Recv nack                          86                        0
Recv duplicated msg-id              5                        0
No TE-link to recv Hop              0                        0
Rcv pkt disabled interface          0                        0
Transmit buffer full                0                        0
Transmit failure                    0                        0
Receive failure                     0                        0
P2MP RESV discarded by appl         0                        0
Rate limit                        306                        0
Err msg loop detected               0                        0
```

show rsvp version

Syntax
show rsvp version
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show rsvp version

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display information about the Resource Reservation Protocol (RSVP) protocol settings, such as the version of the RSVP software, the refresh timer and keep multiplier, and local RSVP graceful restart capabilities on a routing device.

Options
none—Display RSVP protocol settings.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show rsvp version (Router in Steady State) on page 3234
show rsvp version (Router Restarting) on page 3234

Output Fields
Table 433 on page 3233 describes the output fields for the **show rsvp version** command. Output fields are listed in the approximate order in which they appear.

Table 433: show rsvp version Output Fields

Field Name	Field Description
Resource ReSerVation Protocol, version	RSVP software version.
RSVP protocol	Status of RSVP: **Enabled** or **Disabled**.
R(refresh timer)	Configured time interval used to generate periodic RSVP messages.
K(keep multiplier)	Number of RSVP messages that can be lost before an RSVP state is declared stale.
Preemption	Currently configured preemption capability: **Aggressive**, **Disabled**, or **Normal**. The default is **Normal**.
Graceful restart	Status of the graceful restart feature for RSVP on the restarting routing device: **Enabled** or **Disabled**.
Restart helper mode	Status of the helper mode feature: **Enabled** or **Disabled**. When this feature is enabled, the restarting routing device can help the neighbor with its RSVP restart procedures.
Maximum helper restart time	Number of milliseconds (ms) configured for the maximum helper restart time. The maximum helper restart time is the length of time the routing device waits before declaring that an RSVP neighbor attempting to restart gracefully is down.

Table 433: show rsvp version Output Fields *(continued)*

Field Name	Field Description
Maximum helper recovery time	Number of milliseconds configured for the maximum helper recovery time. The maximum helper recovery time is the amount of time the routing device maintains the state of an RSVP neighbor attempting to restart gracefully.
Restart time	Number of milliseconds that a neighbor waits to receive a Hello message from the restarting node before declaring the node dead and deleting the states.
Recovery time	Number of milliseconds during which the restarting node attempts to recover its lost states with help from its neighbors. Recovery time is advertised by the restarting node to its neighbors, and applies to nodal faults. The restarting node considers its graceful restart complete after this time has elapsed.
Soft-preemption cleanup	Time, in seconds, that an LSP is kept after it has been soft preempted. This is a global property of the RSVP protocol.

show rsvp version (Router in Steady State)

```
user@host> show rsvp version
Resource ReSerVation Protocol, version 1. rfc2205
   RSVP protocol             Enabled
   R(refresh timer)          30 seconds
   K(keep multiplier)        3
   Preemption                Normal
   Soft-preemption cleanup   60 seconds
   Graceful restart          Enabled
   Restart helper mode       Enabled
   Restart time              60000 msec
```

show rsvp version (Router Restarting)

```
user@host> show rsvp version
Resource ReSerVation Protocol, version 1. rfc2205
   RSVP protocol:                Enabled
   R(refresh timer):             30 seconds
   K(keep multiplier):           3
   Preemption:                   Normal
   Soft-preemption cleanup:      30 seconds
   Graceful deletion timeout:    30 seconds
   Graceful restart:             Disabled
   Restart helper mode:          Enabled
   Maximum helper restart time:  20000 msec
   Maximum helper recovery time: 180000 msec
   Restart time:                 0 msec
```

show ted database

Syntax
show ted database
<brief | detail | extensive>
<logical-system (all | *logical-system-name*)>
<*system-name*>

Syntax (J-EX Series Switch)
show ted database
<brief | detail | extensive>
<*system-name*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display the entries in the Multiprotocol Label Switching (MPLS) traffic engineering database.

Options
none—Display standard information about all entries in the traffic engineering database.

brief | detail | extensive—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

system-name—(Optional) Display traffic engineering database information for a particular system.

Required Privilege Level
view

List of Sample Output
show ted database brief on page 3237
show ted database detail system-name on page 3238
show ted database extensive on page 3238

Output Fields
Table 434 on page 3235 describes the output fields for the **show ted database** command. Output fields are listed in the approximate order in which they appear.

Table 434: show ted database Output Fields

Field Name	Field Description	Level of Output
TED database	Number of nodes and pseudonodes participating in IS-IS and OSPF domain routing.	All levels
ID	Hostname and address of the node that the link is coming from. An address of **.00** indicates that the node is the routing device itself. An address in the range **0.01** through **0.FF** indicates that the node is a pseudonode. If the node contains a router ID, it is displayed in parentheses.	**brief**
NodeID	Hostname and address of the node that the link is coming from. An address of **.00** indicates that the node is the routing device itself. An address in the range **0.01** through **0.FF** indicates that the node is a pseudonode.	**extensive**
Type	Type of node. It can be either **Rtr** (router) or **Net** (pseudonode).	All levels

Table 434: show ted database Output Fields *(continued)*

Field Name	Field Description	Level of Output
Age(s)	How long since the node was last refreshed, in seconds.	All levels
LnkIn	Number of nodes pointing toward this node.	All levels
LnkOut	Number of nodes to which this node points.	All levels
Protocol	Protocol that reported the node information: • **IS-IS(1)**—IS-IS Level 1. • **IS-IS(2)**—IS-IS Level 2. • **OSPF (area-number)**—OSPF from the specified area.	All levels
To	Address on the far end of a link.	**detail extensive**
Local	Address of the local interface being used to reach the remote node.	**detail extensive**
Remote	Address of the interface on the remote node.	**detail extensive**
Metric	Configured traffic engineering metric.	**extensive**
Static BW	Total interface bandwidth in bps.	**extensive**
Reservable bandwidth	Subscription factor for the interface, which is the percentage of the link bandwidth that can be used for the RSVP reservation process. You configure this by including the **subscription** statement when configuring RSVP.	**extensive**
Available BW [priority]	(Must include **diffserv-te** statement when configuring LSPs) Amount of bandwidth actually reserved by RSVP for each priority level. The bandwidth shown is for the entire interface, not for each individual LSP.	**extensive**
Diffserv-TE BW Model	Bandwidth constraint model used by the LSPs.	**extensive**
Available BW [TE-class]	(Must include the **diffserv-te** statement when configuring LSPs) Amount of bandwidth actually reserved by RSVP for each traffic engineering class.	**extensive**
Static BW [CT-class]	Total interface bandwidth used by an MPLS traffic class, in bps.	**extensive**

Table 434: show ted database Output Fields *(continued)*

Field Name	Field Description	Level of Output

Interface Switching Capability Descriptor (*n*)	Information about the interface switching capability descriptor, which is a subtype length value (TLV) of the link TLV. ***n*** is the index number.
	• **Switching type**—Type of switching to be performed on a particular link:
	  • **PSC-1**—Packet switch-capable 1
	  • **PSC-2**—Packet switch-capable 2
	  • **PSC-3**—Packet switch-capable 3
	  • **PSC-4**—Packet switch-capable 4
	  • **L2SC**—Layer-2-switch-capable
	  • **TDM**—Time-division-multiplexing-capable
	  • **LSC**—Lambda switch-capable
	  • **FSC**—Fiber switch-capable
	• **Encoding type**—Encoding of the LSP being requested:
	  • **Packet**
	  • **Ethernet**
	  • **ANSI/ETSI PDH**
	  • **Reserved**
	  • **SDH /SONET**
	  • **Digital Wrapper**
	  • **Lambda (photonic)**
	  • **Fiber**
	  • **FiberSDH/SONET**
	• **Maximum LSP BW [priority] bps**—Maximum LSP bandwidth information. Amount of bandwidth actually reserved for each priority level. The bandwidth shown is for the entire interface.
	  • **[*n*]**—Priority level. The range is from **0** (high) through **7** (low).
	  • ***n* Mbps**—Amount of the maximum bandwidth.
	• **Minimum LSP BW**—Minimum LSP bandwidth in Mbps. Amount of bandwidth actually reserved for each priority level. The bandwidth shown is for the entire interface. **Minimum LSP BW** is displayed only when **switching type** is **PSC-1** or **TDM**.
	• **Interface MTU**—Displayed only when **switching type** is **TDM**.
	• **Interface supports standard SONET/SDH**—Displayed only when **switching type** is **TDM**.
		extensive

show ted database brief

```
user@host> show ted database brief
TED database: 6 ISIS nodes 6 INET nodes
ID                              Type Age(s) LnkIn LnkOut Protocol
cheviot.00(123.456.1.10)        Rtr    383      1      1 IS-IS(2) IS-IS(1)
corriedale.00(123.456.1.11)     Rtr     36      2      0 IS-IS(2) IS-IS(1)
wolff.00(123.456.1.12)          Rtr    399      0      0 IS-IS(2) IS-IS(1)
perendale.00(123.456.1.13)      Rtr    385      2      0 IS-IS(2) IS-IS(1)
merino.00(123.456.1.14)         Rtr    379      1      3 IS-IS(2) IS-IS(1)
romney.00(123.456.1.15)         Rtr    427      0      2 IS-IS(2) IS-IS(1)
```

show ted database detail system-name

```
user@host> show ted database detail merino
TED database: 6 ISIS nodes 6 INET nodes
NodeID: merino.00(123.456.1.14)
  Type: Rtr, Age: 507 secs, LinkIn: 1, LinkOut: 3
  Protocol: IS-IS(2)
    To: corriedale.00(123.456.1.11), Local: 123.456.8.206, Remote: 123.456.8.207

    To: perendale.00(123.456.1.13), Local: 123.456.8.204, Remote: 123.456.8.205
    To: cheviot.00(123.456.1.10), Local: 123.456.10.65, Remote: 123.456.10.66
  Protocol: IS-IS(1)
    To: corriedale.00(123.456.1.11), Local: 123.456.8.206, Remote: 123.456.8.207

    To: perendale.00(123.456.1.13), Local: 123.456.8.204, Remote: 123.456.8.205
    To: cheviot.00(123.456.1.10), Local: 123.456.10.65, Remote: 123.456.10.66
```

show ted database extensive

```
user@host> show ted database extensive
TED database: 0 ISIS nodes 2 INET nodes
NodeID: 10.255.245.196
  Type: Rtr, Age: 46 secs, LinkIn: 1, LinkOut: 1
  Protocol: OSPF(0.0.0.0)
    To: 10.255.245.24, Local: 4.4.4.4, Remote: 5.5.5.5
      Metric: 1
      Static BW: 155.52Mbps
      Reservable BW: 155.52Mbps
      Available BW [TE-class] bps:
       [te0] 155.52Mbps   [te1] 155.52Mbps   [te2] 155.52Mbps   [te3] 155.52Mbps

       [te4] 155.52Mbps   [te5] 155.52Mbps   [te6] 155.52Mbps   [te7] 155.52Mbps

      Diffserv-TE BW model: Maximum allocation model
        Static BW [CT-class] bps:
         [ct0] 155.52Mbps   [ct1] 155.52Mbps   [ct2] 155.52Mbps   [ct3] 155.52Mbps

        Interface Switching Capability Descriptor(1):
          Switching type: PSC-1
          Encoding type: SDH/SONET
          Maximum LSP BW [priority] bps:
           [0] 155.52Mbps   [1] 155.52Mbps   [2] 155.52Mbps   [3] 155.52Mbps
           [4] 155.52Mbps   [5] 155.52Mbps   [6] 155.52Mbps   [7] 155.52Mbps
          Minimum LSP BW: 155.52Mbps
          Interface MTU: 1285
        Interface Switching Capability Descriptor(2):
          Switching type: TDM
          Encoding type: SDH/SONET
          Maximum LSP BW [priority] bps:
           [0] 155.52Mbps   [1] 155.52Mbps   [2] 155.52Mbps   [3] 155.52Mbps
           [4] 155.52Mbps   [5] 155.52Mbps   [6] 155.52Mbps   [7] 155.52Mbps
          Minimum LSP BW: 155.52Mbps
          Interface supports standard SONET/SDH
```

show ted link

Syntax
show ted link
<brief | detail>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)
show ted link
<brief | detail>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display Multiprotocol Label Switching (MPLS) traffic engineering database link information.

Options
none—Display standard information about traffic engineering database link information.

brief | detail—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show ted link brief on page 3240
show ted link detail on page 3240

Output Fields
Table 435 on page 3239 describes the output fields for the **show ted link** command. Output fields are listed in the approximate order in which they appear.

Table 435: show ted link Output Fields

Field Name	Field Description	Level of Output
ID	Hostname and address of the node that the link is coming from. An address of **.00** indicates that the node is the routing device itself. An address in the range **0.01** through **0.FF** indicates that the node is a pseudonode.	**brief**
-->ID	Hostname and address of the node that the link is going to. An address of **.00** indicates that the node is the routing device itself. An address in the range **0.01** through **0.FF** indicates that the node is a pseudonode.	**brief**
hostname	Hostname and address of the node that the link is coming from. An address of **.00** indicates that the node is the routing device itself. An address in the range **0.01** through **0.FF** indicates that the node is a pseudonode.	**detail**
hostname	Hostname and address of the node that the link is going to. An address of **.00** indicates that the node is the routing device itself. An address in the range **0.01** through **0.FF** indicates that the node is a pseudonode.	**detail**
Local Path	Number of paths CSPF on the local routing device has placed on the link.	All levels
Local BW	Amount of bandwidth the local routing device has placed on the link.	All levels

show ted link brief

```
user@host> show ted link brief
TED link:
ID                          ->ID                          LocalPath LocalBW
cheviot.00(123.456.1.10)      merino.00(123.456.1.14)             0 0bps
merino.00(123.456.1.14)       corriedale.00(123.456.1.11)         0 0bps
merino.00(123.456.1.14)       perendale.00(123.456.1.13)          0 0bps
merino.00(123.456.1.14)       cheviot.00(123.456.1.10)            0 0bps
romney.00(123.456.1.15)       corriedale.00(123.456.1.11)         0 0bps
romney.00(123.456.1.15)       perendale.00(123.456.1.13)          0 0bps
```

show ted link detail

```
user@host> show ted link detail
TED link:
cheviot.00(123.456.1.10)->merino.00(123.456.1.14), LocalPath 0
  localBW [0] 0bps       [1] 0bps      [2] 0bps       [3] 0bps
  localBW [4] 0bps       [5] 0bps      [6] 0bps       [7] 0bps
merino.00(123.456.1.14)->corriedale.00(123.456.1.11), LocalPath 0
  localBW [0] 0bps       [1] 0bps      [2] 0bps       [3] 0bps
  localBW [4] 0bps       [5] 0bps      [6] 0bps       [7] 0bps
merino.00(123.456.1.14)->perendale.00(123.456.1.13), LocalPath 0
  localBW [0] 0bps       [1] 0bps      [2] 0bps       [3] 0bps
  localBW [4] 0bps       [5] 0bps      [6] 0bps       [7] 0bps
merino.00(123.456.1.14)->cheviot.00(123.456.1.10), LocalPath 0
  localBW [0] 0bps       [1] 0bps      [2] 0bps       [3] 0bps
  localBW [4] 0bps       [5] 0bps      [6] 0bps       [7] 0bps
romney.00(123.456.1.15)->corriedale.00(123.456.1.11), LocalPath 0
  localBW [0] 0bps       [1] 0bps      [2] 0bps       [3] 0bps
  localBW [4] 0bps       [5] 0bps      [6] 0bps       [7] 0bps
romney.00(123.456.1.15)->perendale.00(123.456.1.13), LocalPath 0
  localBW [0] 0bps       [1] 0bps      [2] 0bps       [3] 0bps
  localBW [4] 0bps       [5] 0bps      [6] 0bps       [7] 0bps
```

show ted protocol

Syntax

show ted protocol
<brief | detail>
<logical-system (all | *logical-system-name*)>

Syntax (J-EX Series Switch)

show ted protocol
<brief | detail>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display information about the protocols from which the Multiprotocol Label Switching (MPLS) traffic engineering database learned about its nodes.

Options

none—Display standard information about the protocols from which the traffic engineering database learned about its nodes.

brief | detail—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show ted protocol on page 3241

Output Fields

Table 436 on page 3241 describes the output fields for the **show ted protocol** command. Output fields are listed in the approximate order in which they appear.

Table 436: show ted protocol Output Fields

Field Name	Field Description
Protocol name	Protocol that reported the node information: • **IS-IS(1)**—IS-IS Level 1. • **IS-IS(2)**—IS-IS Level 2. • **OSPF** (***area-number***)—OSPF from the specified area.
Credibility	If the protocols provide conflicting information about a node, the protocol with the highest credibility value is the one that the traffic engineering database uses.
Self node	Address the protocol uses as the local address.

show ted protocol

```
user@host> show ted protocol
Protocol name        Credibility  Self node
IS-IS(2)             2  (highest) corriedale.00(123.456.1.11)
IS-IS(1)             1            corriedale.00(123.456.1.11)
```

PART 24

Network Management and Monitoring

CHAPTER 126

Port Mirroring

Port Mirroring—Overview

Understanding Port Mirroring on J-EX Series Switches

Use port mirroring to facilitate analyzing traffic on your J-EX Series Switch on a packet level. Use port mirroring as part of monitoring switch traffic for such purposes as enforcing policies concerning network usage and file sharing, and identifying sources of problems on your network by locating abnormal or heavy bandwidth usage from particular stations or applications.

Port mirroring copies packets to either a local interface for local monitoring or to a VLAN for remote monitoring. You can use port mirroring to copy these packets:

- Packets entering or exiting a port
- Packets entering a VLAN on J-EX4200 Ethernet Switches
- Packets exiting a VLAN on J-EX8200 Ethernet Switches

This topic describes:

Port Mirroring Overview

Port mirroring is needed for traffic analysis on a switch because a switch, unlike a hub, does not broadcast packets to every port on the device. The switch sends packets only to the port to which the destination device is connected. You configure port mirroring on the switch to send copies of unicast traffic to either a local analyzer port or an analyzer

VLAN. Then you can analyze the mirrored traffic using a protocol analyzer application. The protocol analyzer application can run either on a computer connected to the analyzer output interface or on a remote monitoring station.

We recommend that you disable port mirroring when you are not using it and that you select specific interfaces as input to the port mirror analyzer in preference to using the **all** keyword option. You can also limit the amount of mirrored traffic by using statistical sampling, setting a ratio to select a statistical sample, or using a firewall filter. Mirroring only the necessary packets reduces any potential performance impact.

With local port mirroring, traffic from multiple ports is replicated to the analyzer output interface. If the output interface for an analyzer reaches capacity, packets are dropped. You should consider whether the traffic being mirrored exceeds the capacity of the analyzer output interface.

You can use port mirroring on a switch to mirror any of the following:

- **Packets entering or exiting a port**—You can mirror the packets in any combination (on up to 256 ports). For example, you can send copies of the packets entering some ports and the packets exiting other ports to the same local analyzer port or analyzer VLAN.
- **Packets entering a VLAN on a J-EX4200 switch**—You can mirror the packets entering a VLAN on these switches to either a local analyzer port or to an analyzer VLAN. (On J-EX4200 switches, you can configure multiple VLANs [up to 256 VLANs], including a VLAN range and PVLANs, as ingress input to an analyzer.)
- **Packets exiting a VLAN on a J-EX8200 switch**—You can mirror the packets exiting a VLAN on a J-EX8200 switch to either a local analyzer port or to an analyzer VLAN. You can configure multiple VLANs (up to 256 VLANs), including a VLAN range and PVLANs, as egress input to an analyzer.
- **Statistical sample**—You can mirror a statistical sample of packets that are
  - Entering or exiting a port
  - Entering a VLAN on a J-EX4200 switch
  - Exiting a VLAN on a J-EX8200 switch

  You specify the sample number of packets by setting the ratio. You can send the sample to either a local analyzer port or to an analyzer VLAN.
- **Policy-based sample**—You can mirror a policy-based sample of packets that are
  - Entering or exiting a port
  - Entering a VLAN on a J-EX4200 switch
  - Exiting a VLAN on a J-EX8200 switch

  You configure a firewall filter to establish a policy to select certain packets. You can send the sample to a local analyzer port or to an analyzer VLAN.

NOTE: The Junos OS for J-EX Series switches implements port mirroring differently from other Junos OS packages. Junos OS for J-EX Series switches does not include the **port-mirroring** statement found in the **edit forwarding-options** level of the hierarchy of other Junos OS packages, nor the **port-mirror** action in firewall filter terms.

Limitations of Port Mirroring

Port mirroring on J-EX Series switches has the following limitations:

- On a J-EX4200 switch, you can enable only one analyzer (port mirroring configuration).
- On a J-EX8208 or J-EX8216 switch, you can enable a maximum of seven analyzers (port mirroring configurations).
- Packets with physical layer errors are filtered out and thus are not sent to the analyzer port or analyzer VLAN.
- You cannot mirror packets exiting or entering the following ports:
  - Dedicated Virtual Chassis ports (VCPs)
  - Management port (**me0** or **vme0**)
  - Routed VLAN interfaces (RVIs)
- On J-EX8200 switches, you can set a ratio only for ingress packets.
- On J-EX4200 switches, mirrored packets exiting a tagged interface might contain an incorrect VLAN ID.
- On J-EX4200 switches, tagged packets mirrored to an analyzer port might contain an incorrect Ethertype.

Table 437 on page 3247 lists some port mirroring terms and their descriptions.

Port Mirroring Terminology

Table 437: Port Mirroring Terminology

Term	Description
Analyzer	A port-mirroring configuration on a J-EX Series switch. The analyzer includes: • The name of the analyzer • Source (input) ports or VLAN (optional) • A destination for mirrored packets (either a monitor port or a monitor VLAN) • Ratio field for specifying statistical sampling of packets (optional) • Loss-priority setting

Table 437: Port Mirroring Terminology *(continued)*

Term	Description
Analyzer output interface Also known as monitor port	Interface to which mirrored traffic is sent and to which a protocol analyzer application is connected. NOTE: Interfaces used as output for a port mirror analyzer must be configured as family **ethernet-switching**. Analyzer output interfaces have the following limitations: • Cannot also be a source port. • Cannot be used for switching. • Do not participate in Layer 2 protocols, such as Spanning Tree Protocol (STP), when part of a port mirroring configuration. • When configured as an analyzer output interface, they lose any existing VLAN associations. If the bandwidth of the analyzer output interface is not sufficient to handle the traffic from the source ports, overflow packets are dropped.
Analyzer VLAN Also known as monitor VLAN	VLAN to which mirrored traffic is sent. The mirrored traffic can be used by a protocol analyzer application. The monitor VLAN is spread across the switches in your network.
Firewall-based analyzer	An analyzer session that has only an “output” stanza. A firewall-based analyzer must be used along with a firewall filter to achieve the functionality of an analyzer.
Input interface Also known as mirrored ports or monitored interfaces	An interface on the switch that is being mirrored, either on traffic entering or exiting the interface. An input interface cannot also be an output interface for an analyzer.
Mirror ratio	See statistical sampling.
Monitoring station	A computer running a protocol analyzer application.
Native analyzer session	An analyzer session that has both “input” and “output” stanzas.
Policy-based mirroring	Mirroring of packets that match the match items in the defined firewall filter term. The action item **analyzer *analyzer-name*** is used in the firewall filter to send the packets to the port mirror analyzer.
Protocol analyzer application	An application used to examine packets transmitted across a network segment. Also commonly called network analyzer, packet sniffer, or probe.
Remote port mirroring	Functions the same as local port mirroring, except that the mirrored traffic is not copied to a local analyzer port but is flooded into an analyzer VLAN that you create specifically for the purpose of receiving mirrored traffic. In the intermediate switch, you can avoid flooding of the mirrored traffic to the member ports of the VLAN by setting the “ingress only” attribute to the incoming ports of the VLAN and the “egress only” attribute to the outgoing port of the VLAN.

Table 437: Port Mirroring Terminology *(continued)*

Term	Description
Statistical sampling	You can configure the system to mirror a sampling of the packets, by setting a ratio of 1:*x*, where *x* is a value from 1 through 2047. For example, when the ratio is set to 1, all packets are copied to the analyzer. When the ratio is set to 200, 1 of every 200 packets is copied.

Related Documentation

- Example: Configuring Port Mirroring for Local Monitoring of Employee Resource Use on J-EX Series Switches on page 3249
- Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on J-EX Series Switches on page 3254
- Configuring Port Mirroring to Analyze Traffic (J-Web Procedure) on page 3263 or Configuring Port Mirroring to Analyze Traffic (CLI Procedure) on page 3260
- Firewall Filter Match Conditions and Actions for J-EX Series Switches on page 2728

Examples: Port Mirroring Configuration

- Example: Configuring Port Mirroring for Local Monitoring of Employee Resource Use on J-EX Series Switches on page 3249
- Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on J-EX Series Switches on page 3254

Example: Configuring Port Mirroring for Local Monitoring of Employee Resource Use on J-EX Series Switches

J-EX Series switches allow you to configure port mirroring to send copies of packets to either a local interface for local monitoring or to a VLAN for remote monitoring. You can use port mirroring to copy these packets:

- Packets entering or exiting a port
- Packets entering a VLAN on J-EX4200 switches
- Packets exiting a VLAN on J-EX8200 switches

You can analyze the mirrored traffic using a protocol analyzer application installed on a system connected to the local destination interface (or a running on a remote monitoring station if you are sending mirrored traffic to an analyzer VLAN).

This example describes how to configure a J-EX Series switch to mirror traffic entering interfaces connected to employee computers to an analyzer output interface on the same switch.

This example describes how to configure local port mirroring:

- Requirements on page 3250
- Overview and Topology on page 3250

Requirements

This example uses the following hardware and software components:

- One J-EX Series switch

Before you configure port mirroring, be sure you have an understanding of port mirroring concepts.

Overview and Topology

This topic includes two related examples that describe how to mirror traffic entering ports on the switch to a destination interface on the same switch. The first example shows how to mirror all traffic entering the ports connected to employee computers. The second example shows the same scenario, but includes a filter to mirror only the employee traffic going to the Web.

Network Topology

In this example, **ge-0/0/0** and **ge-0/0/1** serve as connections for employee computers.

In this example, one interface, **ge-0/0/10**, is reserved for analysis of mirrored traffic. Connect a PC running a protocol analyzer application to the analyzer output interface to analyze the mirrored traffic.

NOTE: Multiple ports mirrored to one interface can cause buffer overflow and dropped packets.

Figure 84 on page 3250 shows the network topology for this example.

Figure 84: Network Topology for Local Port Mirroring Example

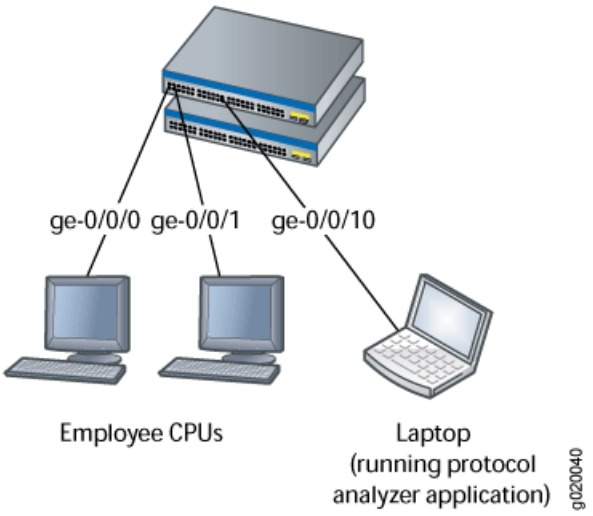

Mirroring All Employee Traffic for Local Analysis

To configure port mirroring for all employee traffic for local analysis, perform these tasks:

CLI Quick Configuration

To quickly configure local port mirroring for ingress traffic to the two ports connected to employee computers, copy the following commands and paste them into the switch terminal window:

```
[edit]
set interfaces ge-0/0/0 unit 0 family ethernet-switching
set interfaces ge-0/0/1 unit 0 family inet 192.1.1.1/24
set interfaces ge-0/0/10 unit 0 family ethernet-switching
set ethernet-switching options analyzer employee–monitor input ingress interface ge-0/0/0.0
set ethernet-switching options analyzer employee–monitor input ingress interface ge-0/0/1.0
set ethernet-switching options analyzer employee–monitor output interface ge-0/0/10.0
```

Step-by-Step Procedure

To configure an analyzer called **employee-monitor** and specify the input (source) interfaces and the analyzer output interface:

1. Configure each interface connected to employee computers as an input interface for the port-mirror analyzer that we are calling **employee-monitor**:

   ```
   [edit ethernet-switching-options]
   user@switch# set analyzer employee-monitor input ingress interface ge–0/0/0.0
   user@switch# set analyzer employee-monitor input ingress interface ge–0/0/1.0
   ```

2. Configure the output analyzer interface for the **employee-monitor** analyzer. This will be the destination interface for the mirrored packets:

   ```
   [edit ethernet-switching-options]
   user@switch# set analyzer employee-monitor output interface ge-0/0/10.0
   ```

Results

Check the results of the configuration:

```
[edit]
user@switch# show
ethernet-switching-options {
    analyzer employee-monitor {
        input {
            ingress {
                interface ge-0/0/0.0;
                interface ge-0/0/1.0;
            }
        }
        output {
            interface {
                ge-0/0/10.0;
            }
        }
    }
}
```

Mirroring Employee-to-Web Traffic for Local Analysis

To configure port mirroring for employee to web traffic, perform these tasks:

CLI Quick Configuration

To quickly configure local port mirroring of traffic from the two ports connected to employee computers, filtering so that only traffic to the external Web is mirrored, copy the following commands and paste them into the switch terminal window:

```
[edit]
set ethernet-switching-options analyzer employee–web–monitor output interface ge-0/0/10.0
set firewall family ethernet-switching filter watch-employee term employee-to-corp from
destination-address 192.0.2.16/28
set firewall family ethernet-switching filter watch-employee term employee-to-corp from
source-address 192.0.2.16/28
set firewall family ethernet-switching filter watch-employee term employee-to-corp then accept
set firewall family ethernet-switching filter watch-employee term employee-to-web from
destination-port 80
set firewall family ethernet-switching filter watch-employee term employee-to-web then analyzer
employee-web-monitor
set interfaces ge-0/0/0 unit 0 family ethernet-switching filter input watch-employee
set interfaces ge-0/0/1 unit 0 family ethernet-switching filter input watch-employee
```

Step-by-Step Procedure

To configure local port mirroring of employee-to-web traffic from the two ports connected to employee computers:

1. Configure the local analyzer interface:

   ```
   [edit interfaces]
   user@switch# set ge-0/0/10 unit 0 family ethernet-switching
   ```

2. Configure the **employee-web-monitor** analyzer output (the input to the analyzer comes from the action of the filter):

   ```
   [edit ethernet-switching-options]
   user@switch# set analyzer employee-web-monitor output interface ge-0/0/10.0
   ```

3. Configure a firewall filter called **watch-employee** to send mirrored copies of employee requests to the Web to the **employee-web-monitor** analyzer. Accept all traffic to and from the corporate subnet (destination or source address of **192.0.2.16/28**). Send mirrored copies of all packets destined for the Internet (**destination port 80**) to the **employee-web-monitor** analyzer.

   ```
   [edit firewall family ethernet-switching]
   user@switch# set filter watch-employee term employee-to-corp from
   destination-address 192.0.2.16/28
   user@switch# set filter watch-employee term employee-to-corp from source-address
   192.0.2.16/28
   user@switch# set filter watch-employee term employee-to-corp then accept
   user@switch# set filter watch-employee term employee-to-web from destination-port
   80
   user@switch# set filter watch-employee term employee-to-web then analyzer
   employee-web-monitor
   ```

4. Apply the **watch-employee** filter to the appropriate ports:

   ```
   [edit interfaces]
   user@switch# set ge-0/0/0 unit 0 family ethernet-switching filter input watch-employee
   user@switch# set ge-0/0/1 unit 0 family ethernet-switching filter input watch-employee
   ```

Results

Check the results of the configuration:

```
[edit]
user@switch# show
```

```
ethernet-switching-options {
   analyzer employee-web-monitor {
      output {
         interface ge-0/0/10.0;
         }
      }
}
...
firewall family ethernet-switching {
   filter watch-employee {
      term employee-to-corp {
         from {
            destination-address 192.0.2.16/28;
            source-address 192.0.2.16/28;
         }
         then accept {
      }
      term employee-to-web {
         from {
            destination-port 80;
         }
         then analyzer employee-web-monitor;
      }
   }
}
...
interfaces {
   ge-0/0/0 {
      unit 0 {
         family ethernet-switching {
            port-mode trunk;
            vlan members [employee-vlan, voice-vlan];
            filter {
               input watch-employee;
            }
         }
      }
   }
   ge-0/0/1 {
      family ethernet-switching {
         filter {
            input watch-employee;
         }
      }
   }
}
```

Verification

To confirm that the configuration is correct, perform these tasks:

- Verifying That the Analyzer Has Been Correctly Created on page 3254

Verifying That the Analyzer Has Been Correctly Created

Purpose Verify that the analyzer named **employee-monitor** or **employee-web-monitor** has been created on the switch with the appropriate input interfaces, and appropriate output interface.

Action You can verify the port mirror analyzer is configured as expected using the **show analyzer** command.

```
user@switch> show analyzer
  Analyzer name                 : employee-monitor
  Output interface              : ge-0/0/10.0
  Mirror ratio                  : 1
  Loss priority                 : Low
  Ingress monitored interfaces : ge-0/0/0.0
  Ingress monitored interfaces : ge-0/0/1.0
  Egress monitored interfaces  : None
```

Meaning This output shows that the **employee-monitor** analyzer has a ratio of 1 (mirroring every packet, the default setting), a loss priority of low (set this option to high only when the analyzer output is to a VLAN), is mirroring the traffic entering the **ge-0/0/0** and **ge-0/0/1** interfaces, and sending the mirrored traffic to the **ge-0/0/10** interface.

Related Documentation

- Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on J-EX Series Switches on page 3254
- Configuring Port Mirroring to Analyze Traffic (CLI Procedure) on page 3260
- Configuring Port Mirroring to Analyze Traffic (J-Web Procedure) on page 3263
- Understanding Port Mirroring on J-EX Series Switches on page 3245

Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on J-EX Series Switches

J-EX Series switches allow you to configure port mirroring to send copies of packets to either a local interface for local monitoring or to a VLAN for remote monitoring. You can use port mirroring to copy these packets:

- Packets entering or exiting a port
- Packets entering a VLAN on J-EX4200 switches
- Packets exiting a VLAN on J-EX8200 switches

You can analyze the mirrored traffic using a protocol analyzer application running on a remote monitoring station if you are sending mirrored traffic to an analyzer VLAN.

This topic includes two related examples that describe how to mirror traffic entering ports on the switch to the **remote-analyzer** VLAN so that you can perform analysis from a remote monitoring station. The first example shows how to mirror all traffic entering the ports connected to employee computers. The second example shows the same scenario, but includes a filter to mirror only the employee traffic going to the Web.

This example describes how to configure remote port mirroring:

Requirements

This example uses the following hardware and software components:

- One J-EX4200 switch connected to another J-EX4200 switch

Before you configure port mirroring, be sure you have an understanding of port mirroring concepts.

Input interfaces that are referred by the analyzer must be configured.

Overview and Topology

This topic includes two related examples that describe how to configure port mirroring to the **remote-analyzer** VLAN so that analysis can be performed from a remote monitoring station. The first example shows how to configure a J-EX Series switch to mirror all traffic from employee computers. The second example shows the same scenario, but the setup includes a filter to mirror only the employee traffic going to the Web.

Figure 85 on page 3255 shows the network topology for this example.

Figure 85: Remote Port Mirroring Example Network Topology

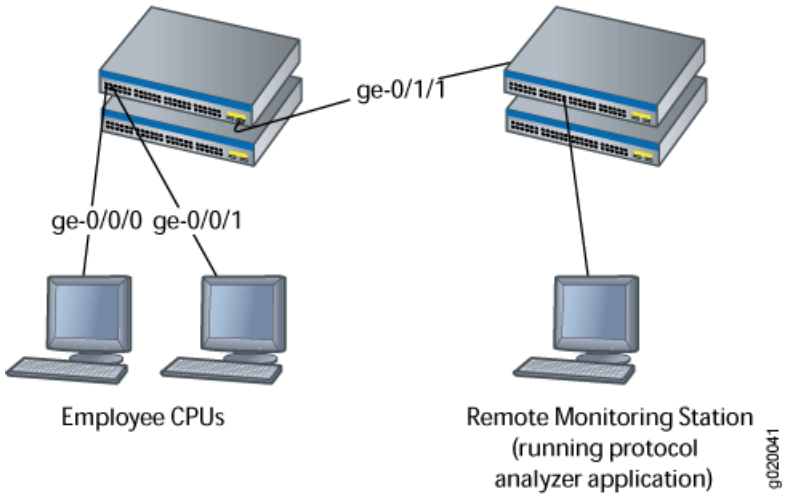

In this example:

- Interface **ge-0/0/0** is a Layer 2 interface and interface **ge-0/0/1** is a Layer 3 interface that serve as connections for employee computers.
- Interface **ge-0/0/10** is a Layer 2 interface that connects to another switch.

- VLAN **remote-analyzer** is configured on all switches in the topology to carry the mirrored traffic.

NOTE: The interface connected to the remote monitoring station must be a member of VLAN remote-analyzer, and this VLAN must be configured on all switches between the monitored switch and the monitoring station.

Mirroring All Employee Traffic for Remote Analysis

To configure port mirroring for remote traffic analysis for all incoming and outgoing employee traffic, perform these tasks:

CLI Quick Configuration

To quickly configure port mirroring for remote traffic analysis for incoming and outgoing employee traffic, copy the following commands and paste them into the switch terminal window:

```
[edit]
set vlans remote-analyzer vlan-id 999
set interfaces ge-0/0/10 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/10 unit 0 family ethernet-switching vlan members 999
set ethernet-switching-options analyzer employee-monitor input ingress interface ge-0/0/0.0
set ethernet-switching-options analyzer employee-monitor input ingress interface ge-0/0/1.0
set ethernet-switching-options analyzer employee-monitor input egress interface ge-0/0/0.0
set ethernet-switching-options analyzer employee-monitor input egress interface ge-0/0/1.0
set ethernet-switching-options analyzer employee-monitor loss-priority high output vlan
remote-analyzer
```

Step-by-Step Procedure

To configure basic remote port mirroring:

1. Configure the VLAN tag ID for the **remote-analyzer** VLAN:

```
[edit vlans]
user@switch# set remote-analyzer vlan-id 999
```

2. Configure the interface on the network port connected to another switch for trunk mode and associate it with the **remote-analyzer** VLAN:

```
[edit interfaces]
user@switch# set ge-0/0/10 unit 0 family ethernet-switching port-mode trunk
user@switch# set ge-0/0/10 unit 0 family ethernet-switching vlan members 999
```

3. Configure the **employee-monitor** analyzer:

```
[edit ethernet-switching-options]
user@switch# set analyzer employee–monitor loss-priority high
user@switch# set analyzer employee-monitor input ingress interface ge-0/0/0.0
user@switch# set analyzer employee–monitor input ingress interface ge-0/0/1.0
user@switch# set analyzer employee–monitor output vlan remote-analyzer
set analyzer employee-monitor input egress interface ge-0/0/0.0
set analyzer employee-monitor input egress interface ge-0/0/1.0
```

Results

Check the results of the configuration:

```
[edit]
user@switch# show
ethernet-switching-options {
```

```
    analyzer employee-monitor {
      loss-priority high;
      input {
        ingress {
          interface ge-0/0/0.0;
          interface ge-0/0/1.0;
        }
        egress {
          interface ge-0/0/0.0;
          interface ge-0/0/1.0;
        }
      }
      output {
        vlan {
          remote-analyzer;
        }
      }
    }
  }
```

Mirroring Employee-to-Web Traffic for Remote Analysis

To configure port mirroring for remote traffic analysis of employee to web traffic, perform these tasks:

CLI Quick Configuration

To quickly configure port mirroring to mirror employee traffic to the external Web, copy the following commands and paste them into the terminal window:

```
[edit]
set ethernet-switching-options analyzer employee-web-monitor loss-priority high output vlan
999
set vlans remote-analyzer vlan-id 999
set interfaces ge-0/0/10 unit 0 family ethernet-switching port mode trunk
set interfaces ge-0/0/10 unit 0 family ethernet-switching vlan members 999
set firewall family ethernet-switching filter watch-employee term employee-to-corp from
destination-address 192.0.2.16/28
set firewall family ethernet-switching filter watch-employee term employee-to-corp from
source-address 192.0.2.16/28
set firewall family ethernet-switching filter watch-employee term employee-to-corp then accept
set firewall family ethernet-switching filter watch-employee term employee-to-web from
destination-port 80
set firewall family ethernet-switching filter watch-employee term employee-to-web then analyzer
employee–web-monitor
set ge-0/0/0 unit 0 family ethernet-switching filter input watch-employee
set interfaces ge-0/0/1 unit 0 family ethernet-switching filter input watch-employee
```

Step-by-Step Procedure

To configure port mirroring of all traffic from the two ports connected to employee computers to the **remote-analyzer** VLAN for use from a remote monitoring station:

1. Configure the **employee-web-monitor** analyzer:

   ```
   [edit ethernet-switching-options]
   user@switch# set interfaces ge-0/0/10 unit 0 family ethernet-switching port mode
   trunk
   user@switch# set analyzer employee-web-monitor loss-priority high output vlan 999
   ```

2. Configure the VLAN tag ID for the **remote-analyzer** VLAN:

   ```
   [edit vlans]
   ```

```
user@switch# set remote-analyzer vlan-id 999
```

3. Configure the interface to associate it with the **remote-analyzer** VLAN:

```
[edit interfaces]
user@switch# set ge-0/0/10 unit 0 family ethernet-switching vlan members 999
```

4. Configure the firewall filter called **watch-employee**:

```
[edit firewall family ethernet-switching]
user@switch# set filter watch-employee term employee-to-corp from
destination-address 192.0.2.16/28
user@switch# set filter watch-employee term employee-to-corp from source-address
192.0.2.16/28
user@switch# set filter watch-employee term employee-to-corp then accept
user@switch# set filter watch-employee term employee-to-web from destination-port
80
user@switch# set filter watch-employee term employee-to-web then analyzer
employee-web-monitor
```

5. Apply the firewall filter to the employee interfaces:

```
[edit interfaces]
user@switch# set ge-0/0/0 unit 0 family ethernet-switching filter input watch-employee
user@switch# set ge-0/0/1 unit 0 family ethernet-switching filter input watch-employee
```

Results Check the results of the configuration:

```
[edit]
user@switch# show
interfaces {
  ...
  ge-0/0/10 {
    unit 0 {
      family ethernet-switching {
        port-mode trunk;
        vlan {
          members remote-analyzer;
        }
      }
    }
  }
  ge-0/0/0 {
    unit 0 {
      family ethernet-switching {
        filter {
          input watch-employee;
        }
      }
    }
  }
  ge-0/0/1 {
    unit 0 {
      family ethernet-switching {
        filter {
          input watch-employee;
        }
      }
```

```
        }
      }
    }
    ...
    firewall {
      family ethernet-switching {
        ...
        filter watch-employee {
          term employee-to-corp {
            from {
              source-address {
                192.0.2.16/28;
              }
              destination-address {
                192.0.2.16/28;
              }
            }
            then accept;
          }
          term employee-to-web {
            from {
              destination-port 80;
            }
            then analyzer employee-web-monitor;
          }
        }
      }
    }
    ethernet-switching-options {
      analyzer employee-web-monitor {
        loss-priority high;
        output {
          vlan {
            999;
          }
        }
      }
    vlans {
      remote-analyzer {
        vlan-id 999;
      }
    }
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying That the Analyzer Has Been Correctly Created on page 3259

Verifying That the Analyzer Has Been Correctly Created

Purpose Verify that the analyzer named **employee-monitor** or **employee-web-monitor** has been created on the switch with the appropriate input interfaces, and appropriate output interface.

Action You can verify the port mirror analyzer is configured as expected using the **show analyzer** command. To view previously created analyzers that are disabled, go to the J-Web interface.

```
user@switch> show analyzer
 Analyzer name                 : employee-monitor
 Output VLAN                   : remote-analyzer
 Mirror ratio                  : 1
 Loss priority                 : High
 Ingress monitored interfaces : ge-0/0/0.0
 Ingress monitored interfaces : ge-0/0/1.0
```

Meaning This output shows that the **employee-monitor** analyzer has a ratio of 1 (mirroring every packet, the default), a loss priority of high (set this option to high whenever the analyzer output is to a VLAN), is mirroring the traffic entering **ge-0/0/0** and **ge-0/0/1**, and sending the mirrored traffic to the analyzer called **remote-analyzer**.

Related Documentation

- Example: Configuring Port Mirroring for Local Monitoring of Employee Resource Use on J-EX Series Switches on page 3249
- Configuring Port Mirroring to Analyze Traffic (CLI Procedure) on page 3260
- Configuring Port Mirroring to Analyze Traffic (J-Web Procedure) on page 3263
- Understanding Port Mirroring on J-EX Series Switches on page 3245

Configuring Port Mirroring

- Configuring Port Mirroring to Analyze Traffic (CLI Procedure) on page 3260
- Configuring Port Mirroring to Analyze Traffic (J-Web Procedure) on page 3263

Configuring Port Mirroring to Analyze Traffic (CLI Procedure)

J-EX Series switches allow you to configure port mirroring to send copies of packets to either a local interface for local monitoring or to a VLAN for remote monitoring. You can use port mirroring to copy these packets:

- Packets entering or exiting a port
- Packets entering a VLAN on J-EX4200 switches
- Packets exiting a VLAN on J-EX8200 switches

We recommend that you disable port mirroring when you are not using it and select specific input interfaces in preference to using the **all** keyword. You can also limit the amount of mirrored traffic by using a firewall filter or the **ratio** keyword to mirror only a selection of packets.

NOTE: If you want to create additional analyzers without deleting the existing analyzer, first disable the existing analyzer using the **disable analyzer** *analyzer-name* command or the J-Web configuration page for port mirroring.

NOTE: Interfaces used as output for a port mirror analyzer must be configured as family ethernet-switching.

- Configuring Port Mirroring for Local Traffic Analysis on page 3261
- Configuring Port Mirroring for Remote Traffic Analysis on page 3261
- Filtering the Traffic Entering an Analyzer on page 3262

Configuring Port Mirroring for Local Traffic Analysis

To mirror interface traffic or VLAN traffic on the switch to an interface on the switch:

1. Choose a name for the port mirroring configuration—in this case, **employee-monitor**—and specify the input—in this case, packets entering **ge-0/0/0** and **ge-0/0/1**:

```
[edit ethernet-switching-options]
user@switch# set analyzer employee-monitor input ingress interface ge–0/0/0.0

[edit ethernet-switching-options]
user@switch# set analyzer employee-monitor input ingress interface ge–0/0/1.0
```

2. Optionally, you can specify a statistical sampling of the packets by setting a ratio:

```
[edit ethernet-switching-options]
user@switch# set analyzer employee-monitor ratio 200
```

   When the ratio is set to 200, 1 of every 200 packets is mirrored to the analyzer. You can use statistical sampling to reduce the volume of mirrored traffic, as a high volume of mirrored traffic can be performance intensive for the switch. On J-EX8200 switches, you can set a ratio only for ingress packets.

3. Configure the destination interface for the mirrored packets:

```
[edit ethernet-switching-options]
user@switch# set analyzer employee-monitor output interface ge-0/0/10.0
```

Configuring Port Mirroring for Remote Traffic Analysis

To mirror traffic that is traversing interfaces or a VLAN on the switch to a VLAN for analysis from a remote location:

1. Configure a VLAN to carry the mirrored traffic. This VLAN is called **remote-analyzer** and given the ID of 999 by convention in this documentation:

```
[edit]
user@switch# set vlans remote-analyzer vlan-id 999
```

2. Set the uplink module interface that is connected to the distribution switch to trunk mode and associate it with the **remote-analyzer** VLAN:

```
[edit]
user@switch# set interfaces ge-0/1/1 unit 0 family ethernet-switching port-mode trunk
vlan members 999
```

3. Configure the analyzer:

   a. Choose a name and set the loss priority to high. Loss priority should always be set to high when configuring for remote port mirroring:

      ```
      [edit ethernet-switching-options]
      user@switch# set analyzer employee–monitor loss-priority high
      ```

   b. Specify the traffic to be mirrored—in this example the packets entering ports **ge-0/0/0** and **ge-0/0/1**:

      ```
      [edit ethernet-switching-options]
      user@switch# set analyzer employee–monitor input ingress interface ge-0/0/0.0

      [edit ethernet-switching-options]
      user@switch# set analyzer employee–monitor input ingress interface ge-0/0/1.0
      ```

   c. Specify the **remote-analyzer** VLAN as the output for the analyzer:

      ```
      [edit ethernet-switching-options]
      user@switch# set analyzer employee-monitor output vlan 999
      ```

4. Optionally, you can specify a statistical sampling of the packets by setting a ratio:

   ```
   [edit ethernet-switching-options]
   user@switch# set analyzer employee-monitor ratio 200
   ```

   When the ratio is set to 200, 1 out of every 200 packets is mirrored to the analyzer. You can use this to reduce the volume of mirrored traffic as a very high volume of mirrored traffic can be performance intensive for the switch.

Filtering the Traffic Entering an Analyzer

To filter which packets are mirrored to an analyzer, create the analyzer and then use it as the action in the firewall filter. You can use firewall filters in both local and remote port mirroring configurations.

If the same analyzer is used in multiple filters or terms, the packets are copied to the analyzer output port or analyzer VLAN only once.

To filter mirrored traffic, create an analyzer and then create a firewall filter. The filter can use any of the available match conditions and must have an action of **analyzer *analyzer-name***. The action of the firewall filter provides the input to the analyzer.

To configure port mirroring with filters:

1. Configure the analyzer name (here, **employee-monitor**) and the output:

   a. For local analysis, set the output to the local interface to which you will connect the computer running the protocol analyzer application:

      ```
      [edit ethernet-switching-options]
      user@switch# set analyzer employee-monitor output interface ge-0/0/10.0
      ```

   b. For remote analysis, set the loss priority to high and set the output to the **remote-analyzer** VLAN:

      ```
      [edit ethernet-switching-options]
      ```

```
user@switch# set analyzer employee—monitor loss-priority high output vlan 999
```

2. Create a firewall filter using any of the available match conditions and specify the action as **analyzer employee-monitor**:

   This step shows a firewall filter called **example-filter**, with two terms:

   a. Create the first term to define the traffic that should not pass through to the analyzer:

   ```
   [edit firewall family ethernet-switching]
   user@switch# set filter example-filter term no-analyzer from source-address ip—address
   ```

   ```
   [edit firewall family ethernet-switching]
   user@switch# set filter example-filter term no-analyzer from destination-address
   ip-address
   ```

   ```
   [edit firewall family ethernet-switching]
   user@switch# set filter example-filter term no-analyzer then accept
   ```

   b. Create the second term to define the traffic that should pass through to the analyzer:

   ```
   [edit firewall family ethernet-switching]
   user@switch# set filter example-filter term to-analyzer from destination-port 80
   ```

   ```
   [edit firewall family ethernet-switching]
   user@switch# set filter example-filter term to-analyzer then analyzer employee—monitor
   ```

3. Apply the firewall filter to the interfaces or VLAN that are input to the analyzer:

   ```
   [edit]
   user@switch# set interfaces ge-0/0/0 unit 0 family ethernet-switching filter input
   example-filter
   ```

   ```
   [edit]
   user@switch# set vlan rspan filter input example-filter
   ```

Related Documentation

- Configuring Port Mirroring to Analyze Traffic (J-Web Procedure) on page 3263
- Example: Configuring Port Mirroring for Local Monitoring of Employee Resource Use on J-EX Series Switches on page 3249
- Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on J-EX Series Switches on page 3254
- Example: Configuring Firewall Filters for Port, VLAN, and Router Traffic on J-EX Series Switches on page 2755
- Understanding Port Mirroring on J-EX Series Switches on page 3245
- Firewall Filters for J-EX Series Switches Overview on page 2721

Configuring Port Mirroring to Analyze Traffic (J-Web Procedure)

J-EX Series switches allow you to configure port mirroring to send copies of packets to either a local interface for local monitoring or to a VLAN for remote monitoring. You can use port mirroring to copy these packets:

- Packets entering or exiting a port
- Packets entering a VLAN on J-EX4200 switches
- Packets exiting a VLAN on J-EX8200 switches

To configure port mirroring on a J-EX Series switch using the J-Web interface:

1. Select **Configure** > **Security** > **Port Mirroring**.

   The first part of the screen displays analyzer details such as the name, status, analyzer port, ratio, and loss priority.

   The second part of the screen lists ingress and egress ports of the selected analyzer.

   NOTE: After you make changes to the configuration in this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options** > **Commit**. See "Using the Commit Options to Commit Configuration Changes (J-Web Procedure)" on page 334 for details about all commit options.

2. Click one:
   - Add—Add an analyzer. Enter information as specified in Table 438 on page 3264.
   - Edit—Modify details of the selected analyzer. Enter information as specified in Table 438 on page 3264.
   - Delete—Delete the selected analyzer.
   - Enable/Disable—Enable or disable the selected analyzer (toggle).

   NOTE: On J-EX4200 switches, only one analyzer can be enabled at a time. On J-EX8200 switches, a maximum of seven analyzers can be enabled.

   NOTE: When an analyzer is deleted or disabled, any filter association is removed.

Table 438: Port Mirroring Configuration Settings

Field	Function	Your Action
Analyzer Name	Specifies the name of the analyzer.	Type a name for the analyzer.

Table 438: Port Mirroring Configuration Settings *(continued)*

Field	Function	Your Action
Ratio	Specifies the ratio of packets to be mirrored. For example: • A ratio of 1 sends copies of all packets. • A ratio of 2047 sends copies of 1 out of every 2047 packets. On J-EX8200 switches, you can set a ratio only for ingress packets.	Enter a number from 0 through 2047.
Loss Priority	Specifies the loss priority of the mirrored packets. By default, the switch applies a lower priority to mirrored data than to regular port-to-port data—mirrored traffic is dropped in preference to regular traffic when capacity is exceeded. For port mirroring configurations with output to an analyzer VLAN, set the loss priority to high.	Keep the default of low, unless the output is to a VLAN.
Analyzer Port	Specifies a local interface or VLAN to which mirrored packets are sent. NOTE: A VLAN must have only one associated interface to be specified as an analyzer interface.	Click **Select**. In the Select Analyzer Port/VLAN window, select either port or VLAN as the **Analyzer Type**. Next, select the required port or VLAN.
Ingress	Specifies interfaces or VLANs for which entering traffic is mirrored.	Click **Add** and select Port or VLAN. Next, select the interfaces or VLANs. Click **Remove** to delete an ingress interface or VLAN.
Egress	Specifies interfaces for which exiting traffic is mirrored.	Click **Add** to add egress interfaces. Click **Remove** to remove egress interfaces.

Related Documentation

- Configuring Port Mirroring to Analyze Traffic (CLI Procedure) on page 3260
- Example: Configuring Port Mirroring for Local Monitoring of Employee Resource Use on J-EX Series Switches on page 3249
- Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on J-EX Series Switches on page 3254
- Understanding Port Mirroring on J-EX Series Switches on page 3245

Verifying Port Mirroring Configuration

- Verifying Input and Output for Port Mirroring Analyzers on J-EX Series Switches on page 3265

Verifying Input and Output for Port Mirroring Analyzers on J-EX Series Switches

Purpose Verify that an analyzer has been created on the switch and has the appropriate output interfaces, and appropriate output interface.

Action You can verify the port mirror analyzer is configured as expected using the **show analyzer** command.

```
[edit]
user@switch>  show analyzer
Analyzer name                  : employee-monitor
  Output VLAN                  : remote-analyzer
  Mirror ratio                 : 1
  Loss priority                : High
  Ingress monitored interfaces : ge-0/0/0.0
  Ingress monitored interfaces : ge-0/0/1.0
```

You can view all of the port mirror analyzers configured on the switch, including any that are disabled, using the **show ethernet-switching-options** command in configuration mode.

```
user@switch# show ethernet-switching-options
inactive: analyzer employee-web-monitor {
    loss-priority high;
    output {

analyzer employee-monitor {
    loss-priority high;
    input {
        ingress {
            interface ge-0/0/0.0;
            interface ge-0/0/1.0;
        }
    }
    output {
        vlan {
            remote-analyzer;
        }
    }
}
```

Meaning This output shows that the employee-monitor analyzer has a ratio of 1 (mirroring every packet, the default), a loss priority of high (set this option to high whenever the analyzer output is to a VLAN), is mirroring the traffic entering **ge-0/0/0** and **ge-0/0/1**, and sending the mirrored traffic to the analyzer called remote-analyzer.

Related Documentation

- Configuring Port Mirroring to Analyze Traffic (J-Web Procedure) on page 3263
- Configuring Port Mirroring to Analyze Traffic (CLI Procedure) on page 3260
- Example: Configuring Port Mirroring for Local Monitoring of Employee Resource Use on J-EX Series Switches on page 3249
- Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on J-EX Series Switches on page 3254
- Understanding Port Mirroring on J-EX Series Switches on page 3245

Configuration Statements for Port Mirroring

- [edit ethernet-switching-options] Configuration Statement Hierarchy on page 3267

[edit ethernet-switching-options] Configuration Statement Hierarchy

```
ethernet-switching-options {
  analyzer {
    name {
      loss-priority priority;
      ratio number;
      input {
        ingress {
          interface (all | interface-name);
          vlan (vlan-id | vlan-name);
        }
        egress {
          interface (all | interface-name);
        }
      output {
        interface interface-name;
        vlan (vlan-id | vlan-name);
      }
    }
  }
  bpdu-block {
    disable-timeout timeout;
    interface (all | [interface-name]);
  }
  dot1q-tunneling {
    ether-type (0x8100 | 0x88a8 | 0x9100);
  }
  interfaces interface-name {
    no-mac-learning;
  }
  mac-notification {
    notification-interval seconds;
  }
  mac-table-aging-time seconds;
  port-error-disable {
    disable-timeout timeout;
  }
  redundant-trunk-group {
    group-name name {
      interface interface-name <primary>;
    }
  }
  secure-access-port {
    dhcp-snooping-file {
      location local_pathname | remote_URL;
      timeout seconds;
      write-interval seconds;
    }
    interface (all | interface-name) {
      allowed-mac {
        mac-address-list;
      }
      (dhcp-trusted | no-dhcp-trusted );
      mac-limit limit action action;
```

```
        no-allowed-mac-log;
        static-ip ip-address {
          vlan vlan-name;
          mac mac-address;
        }
      }
      vlan (all | vlan-name) {
        (arp-inspection | no-arp-inspection );
        dhcp-option82 {
          circuit-id {
            prefix hostname;
            use-interface-description;
            use-vlan-id;
          }
          remote-id {
            prefix hostname | mac | none;
            use-interface-description;
            use-string string;
          }
          vendor-id [string];
        }
        (examine-dhcp | no-examine-dhcp );
        (ip-source-guard | no-ip-source-guard);
        mac-move-limit limit action action;
      }
    }
    storm-control {
      action-shutdown;
      interface (all | interface-name) {
        bandwidth bandwidth;
        no-broadcast;
        no-unknown-unicast;
      }
    }
    traceoptions {
      file filename <files number> <no-stamp> <replace> <size size> <world-readable |
        no-world-readable>;
      flag flag <disable>;
    }
    unknown-unicast-forwarding {
      vlan (all | vlan-name) {
        interface interface-name;
      }
    }
    voip {
      interface (all | [interface-name | access-ports]) {
        vlan vlan-name ;
        forwarding-class <assured-forwarding | best-effort | expedited-forwarding |
          network-control>;
      }
    }
  }
```

Related Documentation

- Understanding Port Mirroring on J-EX Series Switches on page 3245
- Port Security for J-EX Series Switches Overview on page 2545

- Understanding BPDU Protection for STP, RSTP, and MSTP on J-EX Series Switches on page 1278
- Understanding Redundant Trunk Links on J-EX Series Switches on page 1049
- Understanding Storm Control on J-EX Series Switches on page 2511
- Understanding 802.1X and VoIP on J-EX Series Switches on page 2263
- Understanding Q-in-Q Tunneling on J-EX Series Switches on page 1051
- Understanding Unknown Unicast Forwarding on J-EX Series Switches on page 2512
- Understanding MAC Notification on J-EX Series Switches on page 1060

analyzer

Syntax

```
analyzer {
  name {
    ratio number;
    loss-priority priority;
    input {
      ingress {
        interface (all | interface-name);
        vlan (vlan-id | vlan-name);
      }
      egress {
        interface (all | interface-name);
      }
    output {
      interface interface-name;
      vlan (vlan-id | vlan-name);
    }
  }
}
```

Hierarchy Level [edit ethernet-switching-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure port mirroring. One analyzer (port mirroring configuration) can be configured on a J-EX4200 switch and seven analyzers (port mirroring configurations) can be configured on one J-EX8208 or J-EX8216 switch at a time. Other analyzers can be present and disabled.

Default Port mirroring is disabled and the Junos OS creates no default analyzers.

Options *name*—Name that identifies the analyzer. The name can be up to 125 characters long, must begin with a letter, and can include uppercase letters, lowercase letters, numbers, dashes, and underscores. No other special characters are allowed.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Port Mirroring for Local Monitoring of Employee Resource Use on J-EX Series Switches on page 3249
- Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on J-EX Series Switches on page 3254
- Understanding Port Mirroring on J-EX Series Switches on page 3245

egress

Syntax

```
egress {
  interface (all | interface-name);
}
```

Hierarchy Level [edit **ethernet-switching-options analyzer** *name* **input**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify ports for which traffic exiting the interface is mirrored in an port mirroring configuration.

The statement is explained separately.

Default No default.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Understanding Port Mirroring on J-EX Series Switches on page 3245

ethernet-switching-options

Syntax

```
ethernet-switching-options {
  analyzer {
    name {
      loss-priority priority;
      ratio number;
      input {
        ingress {
          interface (all | interface-name);
          vlan (vlan-id | vlan-name);
        }
        egress {
          interface (all | interface-name);
        }
      output {
        interface interface-name;
        vlan (vlan-id | vlan-name);
      }
    }
  }
  bpdu-block {
    disable-timeout timeout;
    interface (all | [interface-name]);
  }
  dot1q-tunneling {
    ether-type (0x8100 | 0x88a8 | 0x9100);
  }
  interfaces interface-name {
    no-mac-learning;
  }
  mac-notification {
    notification-interval seconds;
  }
  mac-table-aging-time seconds;
  port-error-disable {
    disable-timeout timeout;
  }
  redundant-trunk-group {
    group-name name {
      interface interface-name <primary>;
      interface interface-name;
    }
  }
  secure-access-port {
    dhcp-snooping-file {
      location local_pathname | remote_URL;
      timeout seconds;
      write-interval seconds;
    }
    interface (all | interface-name) {
      allowed-mac {
        mac-address-list;
      }
```

```
        (dhcp-trusted | no-dhcp-trusted);
        mac-limit limit action action;
        no-allowed-mac-log;
        static-ip ip-address {
          vlan  vlan-name;
          mac mac-address;
        }
      }
      vlan (all | vlan-name) {
        (arp-inspection | no-arp-inspection);
        dhcp-option82 {
          circuit-id {
            prefix hostname;
            use-interface-description;
            use-vlan-id;
          }
          remote-id {
            prefix hostname | mac | none;
            use-interface-description;
            use-string string;
          }
          vendor-id [string];
        }
        (examine-dhcp | no-examine-dhcp);
        (ip-source-guard | no-ip-source-guard);
        mac-move-limit limit action action;
      }
    }
    storm-control {
      action-shutdown;
      interface (all | interface-name) {
        bandwidth bandwidth;
        no-broadcast;
        no-unknown-unicast;
      }
    }
    traceoptions {
      file filename <files number> <no-stamp> <replace> <size size> <world-readable |
        no-world-readable>;
      flag flag <disable>;
    }
    unknown-unicast-forwarding {
      vlan (all | vlan-name) {
        interface interface-name;
      }
    }
    voip {
      interface (all | [interface-name | access-ports]) {
        vlan vlan-name ;
        forwarding-class <assured-forwarding | best-effort | expedited-forwarding |
          network-control>;
      }
    }
  }
```

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure Ethernet switching options.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing–control—To add this statement to the configuration.

Related Documentation

- Understanding Port Mirroring on J-EX Series Switches on page 3245
- Port Security for J-EX Series Switches Overview on page 2545
- Understanding BPDU Protection for STP, RSTP, and MSTP on J-EX Series Switches on page 1278
- Understanding Redundant Trunk Links on J-EX Series Switches on page 1049
- Understanding Storm Control on J-EX Series Switches on page 2511
- Understanding 802.1X and VoIP on J-EX Series Switches on page 2263
- Understanding Q-in-Q Tunneling on J-EX Series Switches on page 1051
- Understanding Unknown Unicast Forwarding on J-EX Series Switches on page 2512
- Understanding MAC Notification on J-EX Series Switches on page 1060

ingress

Syntax

```
ingress {
  interface (all | interface-name);
  vlan (vlan-id | vlan-name);
}
```

Hierarchy Level [edit ethernet-switching-options analyzer *name* input]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure ports or VLANs for which the entering traffic is mirrored as part of an port mirroring configuration.

The statements are explained separately.

Default No default.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Port Mirroring for Local Monitoring of Employee Resource Use on J-EX Series Switches on page 3249
- Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on J-EX Series Switches on page 3254
- Understanding Port Mirroring on J-EX Series Switches on page 3245

input

Syntax

```
input {
  ingress {
    interface (all | interface-name);
    vlan (vlan-id | vlan-name);
  }
  egress {
    interface (all | interface-name);
  }
}
```

Hierarchy Level [edit **ethernet-switching-options analyzer** *name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define the traffic to be mirrored in a port mirroring configuration—the definition can be a combination of:

- Packets entering or exiting a port
- Packets entering a VLAN on a J-EX4200 switch
- Packets exiting a VLAN on a J-EX8200 switch

The remaining statements are explained separately.

Default No default.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Port Mirroring for Local Monitoring of Employee Resource Use on J-EX Series Switches on page 3249
- Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on J-EX Series Switches on page 3254
- Understanding Port Mirroring on J-EX Series Switches on page 3245

interface

Syntax interface (all | *interface-name*);

Hierarchy Level [edit **ethernet-switching-options analyzer** *name* **input egress**],
[edit **ethernet-switching-options analyzer** *name* **input ingress**],
[edit **ethernet-switching-options analyzer** *name* **output**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the interfaces for which traffic is mirrored.

Options all—Apply port mirroring to all interfaces on the switch. Mirroring a high volume of traffic can be performance intensive for the switch. Therefore, you should generally select specific input interfaces in preference to using the all keyword, or use the all keyword in combination with setting a ratio for statistical sampling.

interface-name—Apply port mirroring to the specified interface only.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Port Mirroring for Local Monitoring of Employee Resource Use on J-EX Series Switches on page 3249
- Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on J-EX Series Switches on page 3254
- Understanding Port Mirroring on J-EX Series Switches on page 3245

loss-priority

Syntax loss-priority *priority*;

Hierarchy Level [edit **ethernet-switching-options analyzer** *name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a loss priority for mirrored packets. By default, the switch applies a lower priority to mirrored data than to regular port-to-port data—mirrored traffic is dropped in preference for regular traffic when capacity is exceeded. For port mirroring configurations with output to an analyzer VLAN, set the loss priority to high.

Default Low

Options *priority*—The value for priority can be low or high.

Default: low

Required Privilege Level routing—To view this statement in the configuration.routing-control—To add this statement to the configuration.

Related Documentation

- Understanding Port Mirroring on J-EX Series Switches on page 3245
- Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on J-EX Series Switches on page 3254

output

Syntax

```
output {
  interface interface-name;
  vlan (vlan-id | vlan-name);
}
```

Hierarchy Level [edit ethernet-switching-options analyzer *name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the destination for mirrored traffic, either an interface on the switch, for local monitoring, or a VLAN, for remote monitoring.

The statements are explained separately.

Default No default.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Port Mirroring for Local Monitoring of Employee Resource Use on J-EX Series Switches on page 3249
- Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on J-EX Series Switches on page 3254
- Understanding Port Mirroring on J-EX Series Switches on page 3245

ratio

Syntax ratio *number*;

Hierarchy Level [edit **ethernet-switching-options analyzer** *name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure port mirroring to copy a sampling of packets, by setting a ratio of 1:x, A ratio of 1 mirrors all packets, and 2047 mirrors 1 out of every 2047 packets.

On J-EX8200 switches, you can set a ratio only for ingress packets.

Default 1

Options ***number***—The number of packets in the sample, out of which 1 packet is mirrored.
Range: 1 through 2047
Default: 1

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Understanding Port Mirroring on J-EX Series Switches on page 3245

vlan

Syntax vlan (*vlan-id* | *vlan-name*);

Hierarchy Level [edit **ethernet-switching-options analyzer** *name* **input ingress**],
[edit **ethernet-switching-options analyzer** *name* **output**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure mirrored traffic to be sent to a VLAN for remote monitoring.

Options *vlan-id*—Numeric VLAN identifer.

vlan-name—Name of the VLAN.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring Port Mirroring for Remote Monitoring of Employee Resource Use on J-EX Series Switches on page 3254
- Understanding Port Mirroring on J-EX Series Switches on page 3245

Operational Mode Commands for Port Mirroring

show analyzer

Syntax **show analyzer *analyzer-name***

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display information about analyzers configured for port mirroring.

Options *analyzer-name*—(Optional) Displays the status of a specific analyzer on the switch.

Required Privilege Level view

List of Sample Output **show analyzer on page 3281**

Output Fields Table 439 on page 3281 lists the output fields for the **command-name** command. Output fields are listed in the approximate order in which they appear.

Table 439: show analyzer Output Fields

Field Name	Field Description
Analyzer name	Displays the name of the analyzer.
Output interface	Specifies a local interface to which mirrored packets are sent. An analyzer can have output to either an interface or a VLAN, not both.
Output VLAN	Specifies a VLAN to which mirrored packets are sent. An analyzer can have output to either an interface or a VLAN, not both.
Mirror ratio	Displays the ratio of packets to be mirrored, between 1 and 2047 where 1 sends copies of all packets and 2047 sends copies of 1 out of every 2047 packets.
Loss priority	Displays the loss priority of mirrored packets. By default, loss priority is set to **low**, with mirrored traffic dropped in preference for regular traffic when capacity is exceeded. For analyzers with output to a VLAN, set the loss priority to **high**.
Egress monitored interfaces	Displays interfaces for which traffic exiting the interfaces is mirrored.
Ingress monitored interfaces	Displays interfaces for which traffic entering the interfaces is mirrored.
Ingress monitored VLANs	Displays VLANs for which traffic entering the VLAN is mirrored.

show analyzer

```
user@host> show analyzer
Analyzer name                  : employee-monitor
  Output interface             : ge-0/0/10.0
  Output VLAN                  : remote-analyzer
  Mirror ratio                 : 1
  Loss priority                : High
  Egress monitored interfaces  : ge-0/0/3.0
  Ingress monitored interfaces : ge-0/0/0.0
```

```
Ingress monitored interfaces : ge-0/0/1.0
```

CHAPTER 127

sFlow Monitoring Technology

sFlow Technology—Overview

Understanding How to Use sFlow Technology for Network Monitoring on a J-EX Series Switch

The sFlow technology is a monitoring technology for high-speed switched or routed networks. sFlow monitoring technology randomly samples network packets and sends the samples to a monitoring station. You can configure sFlow technology on a J-EX Series Switch to continuously monitor traffic at wire speed on all interfaces simultaneously.

This topic describes:

Sampling Mechanism and Architecture of sFlow Technology on J-EX Series Switches

sFlow technology uses the following two sampling mechanisms:

- Packet-based sampling: Samples one packet out of a specified number of packets from an interface enabled for sFlow technology.
- Time-based sampling: Samples interface statistics at a specified interval from an interface enabled for sFlow technology.

The sampling information is used to create a network traffic visibility picture. The Junos OS fully supports the sFlow standard described in RFC 3176, *InMon Corporation's sFlow:*

A Method for Monitoring Traffic in Switched and Routed Networks (see http://faqs.org/rfcs/rfc3176.html).

NOTE: sFlow technology on the switches samples only raw packet headers. A raw Ethernet packet is the complete Layer 2 network frame.

An sFlow monitoring system consists of an sFlow agent embedded in the switch and a centralized collector. The sFlow agent's two main activities are random sampling and statistics gathering. It combines interface counters and flow samples and sends them across the network to the sFlow collector.

J-EX Series switches adopt the distributed sFlow architecture. The sFlow agent has two separate sampling entities that are associated with each Packet Forwarding Engine. These sampling entities are known as subagents. Each subagent has a unique ID that is used by the collector to identify the data source. A subagent has its own independent state and forwards its own sample messages to the sFlow agent. The sFlow agent is responsible for packaging the samples into datagrams and sending them to the sFlow collector. Because sampling is distributed across subagents, the protocol overhead associated with sFlow technology is significantly reduced at the collector. If the mastership assignment changes in a Virtual Chassis setup, sFlow technology continues to function.

Adaptive Sampling

The switches use adaptive sampling to ensure both sampling accuracy and efficiency. Adaptive sampling is a process of monitoring the overall incoming traffic rate on the network device and providing intelligent feedback to interfaces to dynamically adapt their sampling rate to the traffic conditions. Interfaces on which incoming traffic exceeds the system threshold are checked so that all violations can be regulated without affecting the traffic on other interfaces. Every 5 seconds the agent checks interfaces to get the number of samples, and interfaces are grouped based on the slot that they belong to. The top five interfaces that produce the highest number of samples are selected. Using the binary backoff algorithm, the sampling load on these interfaces is reduced by half and allotted to interfaces that have a lower sampling rate. Therefore when the processor limit is reached, the sampling rate is adapted such that it does not load the processor any further. If the switch is rebooted, the adaptive sampling rate is reset to the user-configured sampling rate. Also, if you modify the sampling rate, the adaptive sampling rate changes.

The advantage of adaptive sampling is that the switch continues to operate at its optimum level even when there is a change in the traffic patterns in the interfaces. You do not need to make any changes. Because the sampling rate adapts dynamically to changing network conditions, the resources are utilized optimally resulting in a high performance network.

Infrequent sampling flows are not reported in the sFlow information, but over time the majority of flows are reported. Based on a defined sampling rate, 1 out of *N* packets is captured and sent to the collector. This type of sampling does not provide a 100 percent accurate result in the analysis, but it does provide a result with quantifiable accuracy. A polling interval defines how often the sFlow data for a specific interface are sent to the collector, but an sFlow agent can also schedule polling.

NOTE: sFlow technology on J-EX Series switches does not support graceful restart. When a graceful restart occurs, the adaptive sampling rate is set to the user-configured sampling rate.

sFlow Agent Address Assignment

The sFlow collector uses the sFlow agent's IP address to determine the source of the sFlow data. You can configure the IP address of the sFlow agent to ensure that the agent ID for the sFlow agent remains constant. If you do not specify the IP address to be assigned to the agent, the IP address assigned to the agent is based on the following order of priority of interfaces configured on the switch:

1. Virtual management Ethernet (VME) interface

2. Management Ethernet interface

If neither of the preceding interfaces has been configured, the IP address of any Layer 3 interface or the routed VLAN interface (RVI) is used as the IP address for the agent. At least one interface must be configured on the switch for an IP address to be automatically assigned to the agent. When the agent IP address is assigned automatically, the IP address is dynamic and changes when the switch reboots.

sFlow data can be used to provide network traffic visibility information. You can explicitly configure the IP address to be assigned to source data (sFlow datagrams). If you do not explicitly configure that address, the IP address of the configured Gigabit Ethernet interface, 10-Gigabit Ethernet interface, or the routed VLAN interface (RVI) is used as the source IP address.

Related Documentation

- Example: Configuring sFlow Technology to Monitor Network Traffic on J-EX Series Switches on page 3285
- Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 3290
- Monitoring Interface Status and Traffic on page 931

Example: sFlow Technology Configuration

- Example: Configuring sFlow Technology to Monitor Network Traffic on J-EX Series Switches on page 3285

Example: Configuring sFlow Technology to Monitor Network Traffic on J-EX Series Switches

You can configure sFlow technology, designed for monitoring high-speed switched or routed networks, to continuously monitor traffic at wire speed on all interfaces simultaneously. sFlow data can be used to provide network traffic visibility information.

This example describes how to configure and use sFlow monitoring. The Junos OS fully supports the sFlow standard described in RFC 3176, *InMon Corporation's sFlow: A Method*

for Monitoring Traffic in Switched and Routed Networks (see http://faqs.org/rfcs/rfc3176.html).

Requirements

This example uses the following hardware and software components:

- One J-EX Series switch

Overview and Topology

sFlow technology is a statistical-sampling–based network monitoring technology for high-speed switched or routed networks. sFlow technology samples network packets and sends the samples to a monitoring station. The information gathered is used to create a network traffic visibility picture.

An sFlow monitoring system consists of an sFlow agent embedded in the switch and a centralized collector. The sFlow agent runs on the switch. It combines interface counters and flow samples and sends them across the network to the sFlow collector. Figure 86 on page 3287 depicts the basic elements of the sFlow system.

Figure 86: sFlow Technology Monitoring System

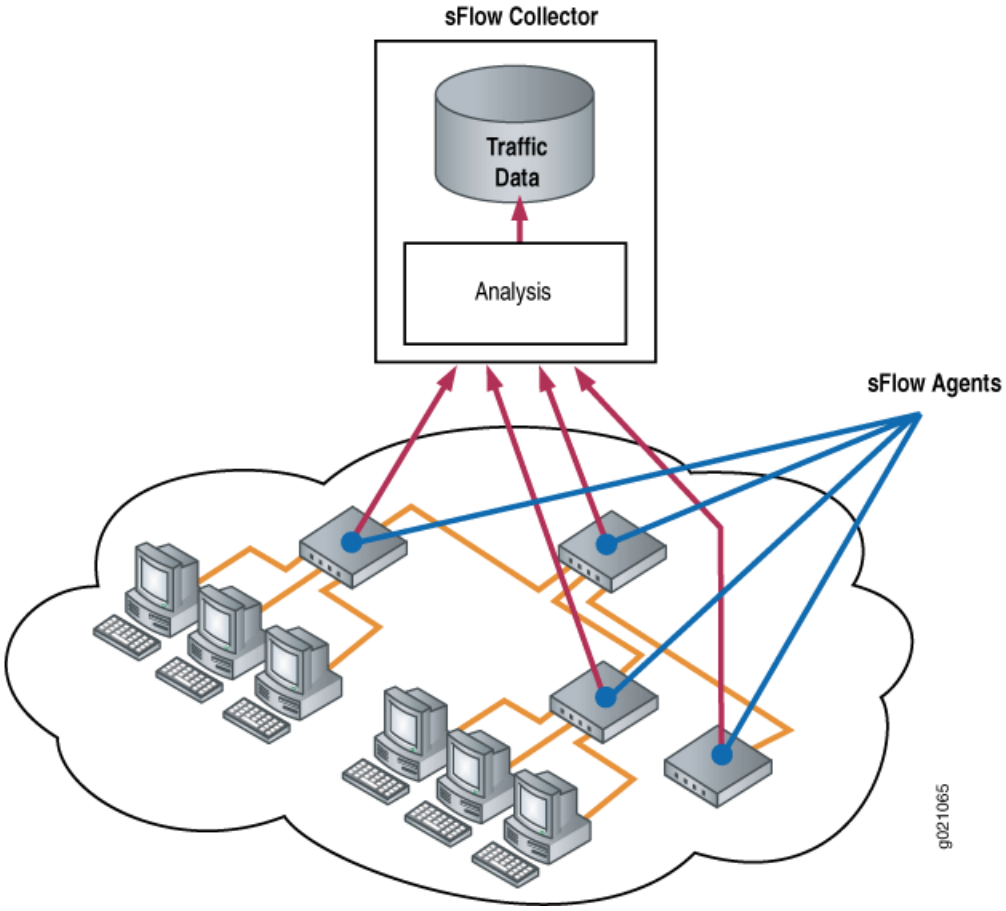

Configuration

To configure sFlow technology, perform the following tasks:

CLI Quick Configuration

To quickly configure sFlow technology, copy the following commands and paste them into the switch terminal window:

```
[edit protocols sflow]
set collector 10.204.32.46
Set collector udp-port 5600
set interfaces ge-0/0/0
set polling-interval 20
set sample-rate 1000
```

Step-by-Step Procedure

To configure sFlow technology:

1. Configure the IP address of the collector:

```
[edit protocols sflow]
user@switch# set collector 10.204.32.46
```

NOTE: You can configure a maximum of 4 collectors.

2. Configure the UDP port of the collector. The default UDP port assigned is 6343.

```
[edit protocols sflow]
user@switch# set collector udp-port 5600
```

3. Enable sFlow technology on a specific interface:

```
[edit protocols sflow]
user@switch# set interfaces ge-0/0/0
```

NOTE: You cannot enable sFlow technology on a Layer 3 VLAN-tagged interface.

You cannot enable sFlow technology on a link aggregation group (LAG) interface--that is, an aggregated Ethernet interface with a name such as ae0. You can enable sFlow technology on the member interfaces that make up the LAG.

4. Specify how often the sFlow agent polls the interface:

```
[edit protocols sflow]
user@switch# set polling-interval 20
```

NOTE: The polling interval can be specified as a global parameter also. Specify 0 if you do not want to poll the interface.

5. Specify the rate at which packets must be sampled:

```
[edit protocols sflow]
user@switch# set sample-rate 1000
```

Results

Check the results of the configuration:

```
[edit protocols sflow]
  user@switch# show
  polling-interval 20;
  udp-port 5600;
  sample-rate 1000;
  collector 10.204.32.46;
  interfaces ge-0/0/0.0;
```

Verification

To confirm that the configuration is correct, perform these tasks:

- Verifying That sFlow Technology Has Been Configured Properly on page 3289
- Verifying That sFlow Technology Is Enabled on the Intended Interface on page 3289
- Verifying the sFlow Collector Configuration on page 3290

Verifying That sFlow Technology Has Been Configured Properly

Purpose Verify that sFlow technology has been configured properly.

Action Use the **show sflow** command:

```
user@switch> show sflow
sFlow     : Enabled
Sample limit  : 300 packets/second
Polling interval : 20 seconds
Sample rate   : 1:1000
Agent ID    : 10.204.96.222
```

NOTE: The sample limit cannot be configured and is set to 300 packets/second.

Meaning The output shows that sFlow technology is enabled and specifies the values for the sample rate, sample limit, and polling interval.

Verifying That sFlow Technology Is Enabled on the Intended Interface

Purpose Verify that sFlow technology is enabled on interfaces and display the sampling parameters.

Action Use the **show sflow interface** command:

```
user@switch> show sflow interface
Interface   Status  Sample rate  Adapted sample rate  Polling-interval
ge-0/0/0.0  Enabled    1000       16000                20
```

NOTE: The sample limit cannot be configured and is set to 300 packets/second.

Meaning The output indicates that sFlow technology is enabled on the **ge-0/0/0.0** interface with a sampling rate of 1000, sampling limit of 300 packets per second and a polling interval of 20 seconds.

Verifying the sFlow Collector Configuration

Purpose Verify the sFlow collector's configuration.

Action Use the **show sflow collector** command:

```
user@switch> show sflow collector
Collector address  UDP-port  No of samples
10.204.32.46       5600      1000
10.204.32.76       3400      1000
```

Meaning The output displays the IP address of the collectors and the UDP ports. It also displays the number of samples.

Related Documentation

- Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 3290
- Understanding How to Use sFlow Technology for Network Monitoring on a J-EX Series Switch on page 3283

Configuring sFlow Technology

Configuring sFlow Technology for Network Monitoring (CLI Procedure)

You can configure sFlow technology, designed for monitoring high-speed switched or routed networks, to continuously monitor traffic at wire speed on all interfaces simultaneously. The Junos OS fully supports the sFlow standard described in RFC 3176, *InMon Corporation's sFlow: A Method for Monitoring Traffic in Switched and Routed Networks* (see http://faqs.org/rfcs/rfc3176.html).`

To configure sFlow features:

1. Configure the IP address of the collector:

   ```
   [edit protocols sflow]
   user@switch# set collector ip-address
   ```

2. Configure the UDP port of the collector. The default UDP port assigned is 6343.

   ```
   [edit protocols sflow]
   user@switch# set collector udp-port port-number
   ```

3. Enable sFlow technology on a specific interface:

   ```
   [edit protocols sflow]
   user@switch# set interfaces interface-name
   ```

NOTE: You cannot enable sFlow technology on a Layer 3 VLAN-tagged interface.

You cannot enable sFlow technology on a link aggregation group (LAG), but you can enable it on the member interfaces of a LAG.

4. Specify how often the sFlow agent polls the interface:

```
[edit protocols sflow]
user@switch# set polling-interval seconds
```

NOTE: Specify 0 if you do not want to poll the interface.

5. Specify the rate at which packets must be sampled:

```
[edit protocols sflow]
user@switch# set sample-rate number
```

6. To configure the polling interval and sample rate at the interface level:

```
[edit protocols sflow interfaces interface-name]
user@switch# set polling-interval seconds

[edit protocols sflow interfaces]
user@switch# set sample-rate number
```

NOTE: The interface-level configuration overrides the global configuration.

7. To specify an IP address to be used as the agent ID for the sFlow agent:

```
[edit protocols sflow]
user@switch# set agent-id ip-address
```

8. To specify the source IP address to be used for sFlow datagrams:

```
[edit protocols sflow]
user@switch# set source-ip ip-address
```

Related Documentation

- Example: Monitoring Network Traffic Using sFlow Technology on J-EX Series Switches on page 3285
- Understanding How to Use sFlow Technology for Network Monitoring on a J-EX Series Switch on page 3283

Configuration Statements for sFlow Technology

- [edit protocols] Configuration Statement Hierarchy on page 3292

[edit protocols] Configuration Statement Hierarchy

```
protocols {
  connections {
    remote-interface-switch connection-name {
      interface interface-name.unit-number;
      transmit-lsp label-switched-path;
      receive-lsp label-switched-path;
    }
  }
  dot1x {
    authenticator {
      authentication-profile-name profile-name;
      interface (all | [ interface-names ]) {
        disable;
        guest-vlan ( vlan-id | vlan-name);
        mac-radius <restrict>;
        maximum-requests number;
        no-reauthentication;
        quiet-period seconds;
        reauthentication {
          interval seconds;
        }
        retries number;
        server-fail (deny | permit | use-cache | vlan-id | vlan-name);
        server-reject-vlan (vlan-id | vlan-name);
        server-timeout seconds;
        supplicant (multiple | single | single-secure);
        supplicant-timeout seconds;
        transmit-period seconds;
      }
      static mac-address {
        interface interface-name;
        vlan-assignment (vlan-id |vlan-name);
      }
    }
  }
  gvrp {
    <enable | disable>;
    interface (all | [interface-name]) {
      disable;
    }
    join-timer millseconds;
    leave-timer milliseconds;
    leaveall-timer milliseconds;
  }
  igmp-snooping {
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>
        <match regex>;
      flag flag (detail | disable | receive | send);
    }
    vlan (vlan-id | vlan-number) {
      data-forwarding {
        source {
          groups group-prefix;
```

```
        }
        receiver {
          source-vlans vlan-list;
          install ;
        }
      }
      disable {
        interface interface-name
      }
      immediate-leave;
      interface interface-name {
        group-limit limit;
        multicast-router-interface;
        static {
          group ip-address;
        }
      }
      proxy ;
      query-interval seconds;
      query-last-member-interval seconds;
      query-response-interval seconds;
      robust-count number;
    }
  }
  lldp {
    disable;
    advertisement-interval seconds;
    hold-multiplier number;
    interface (all | interface-name) {
      disable;
    }
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>
        <match regex>;
      flag flag (detail | disable | receive | send);
    }
  }
  lldp-med {
    disable;
    fast-start number;
    interface (all | interface-name) {
      disable;
      location {
        elin number;
        civic-based {
          what number;
          country-code code;
          ca-type {
            number {
              ca-value value;
            }
          }
        }
      }
    }
  }
```

```
mpls {
  interface ( all | interface-name );
  label-switched-path lsp-name to remote-provider-edge-switch;
  path destination {
    <address | hostname> <strict | loose>
  }
mstp {
  disable;
  bpdu-block-on-edge;
  bridge-priority priority;
  configuration-name name;
  forward-delay seconds;
  hello-time seconds;
  interface (all | interface-name) {
    disable;
    bpdu-timeout-action {
      block;
      alarm;
    }
    cost  cost;
    edge;
    mode mode;
    no-root-port;
    priority priority;
  }
  max-age seconds;
  max-hops hops;
  msti msti-id {
    vlan (vlan-id | vlan-name);
    interface interface-name {
      disable;
      cost  cost;
      edge;
      mode mode;
      priority priority;
    }
  }
  revision-level revision-level;
  traceoptions {
    file filename <files number > <size size> <no-stamp | world-readable |
      no-world-readable>;
    flag flag;
  }
}
mvrp {
  disable
  interface (all | interface-name) {
    disable;
    join-timer milliseconds;
    leave-timer milliseconds;
    leaveall-timer milliseconds;
    registration (forbidden | normal);
  }
  no-dynamic-vlan;
  traceoptions {
```

```
    file filename <files number > <size size> <no-stamp | world-readable |
      no-world-readable>;
    flag flag;
  }
}
oam {
  ethernet{
    connectivity-fault-management {
      action-profile profile-name {
        default-actions {
          interface-down;
        }
      }
      linktrace {
        age (30m | 10m | 1m | 30s | 10s);
        path-database-size path-database-size;
      }
      maintenance-domain domain-name {
        level number;
        mip-half-function (none | default |explicit);
        name-format (character-string | none | dns | mac+2oct);
        maintenance-association ma-name {
          continuity-check {
            hold-interval minutes;
            interval (10m | 10s | 1m | 1s| 100ms);
            loss-threshold number;
          }
          mep mep-id {
            auto-discovery;
            direction down;
            interface interface-name;
            remote-mep mep-id {
              action-profile profile-name;
            }
          }
        }
      }
    }
    link-fault-management {
      action-profile profile-name;
        action {
          syslog;
          link-down;
        }
        event {
          link-adjacency-loss;
          link-event-rate;
          frame-error count;
          frame-period count;
          frame-period-summary count;
          symbol-period count;
        }
        interface interface-name {
          link-discovery (active | passive);
          pdu-interval interval;
          event-thresholds threshold-value;
```

```
            remote-loopback;
            event-thresholds {
              frame-errorcount;
              frame-period count;
              frame-period-summary count;
              symbol-period count;
            }
          }
          negotiation-options {
            allow-remote-loopback;
            no-allow-link-events;
          }
        }
      }
    }
    rstp {
      disable;
      bpdu-block-on-edge;
      bridge-priority priority;
      forward-delay seconds;
      hello-time seconds;
      interface (all | interface-name) {
        disable;
        bpdu-timeout-action {
          block;
          alarm;
        }
        cost  cost;
        edge;
        mode mode;
        no-root-port;
        priority priority;
      }
      max-age seconds;
      }
      traceoptions {
        file filename <files number > <size size> <no-stamp | world-readable |
          no-world-readable>;
        flag flag;
      }
    }
    sflow {
      agent-id
      collector {
        ip-address;
        udp-port port-number;
      }
      disable;
      interfaces interface-name {
        disable;
        polling-interval seconds;
        sample-rate number;
      }
      polling-interval seconds;
      sample-rate number;
      source-ip
```

```
    }
    stp {
      disable;
      bridge-priority priority;
      forward-delay seconds;
      hello-time seconds;
      interface (all | interface-name) {
        disable;
        bpdu-timeout-action {
          block;
          alarm;
        }
        cost  cost;
        edge;
        mode mode;
        no-root-port;
        priority priority;
      }
      max-age seconds;
    }
      traceoptions {
        file filename <files number > <size size> <no-stamp | world-readable |
          no-world-readable>;
        flag flag;
      }
    vstp {
      bpdu-block-on-edge;
      disable;
      force-version stp;
      vlan (all | vlan-id | vlan-name) {
        bridge-priority priority;
        forward-delay seconds;
        hello-time seconds;
        interface (all | interface-name) {
          bpdu-timeout-action {
            alarm;
            block;
          }
          cost  cost;
          disable;
          edge;
          mode mode;
          no-root-port;
          priority priority;
        }
        max-age seconds;
        traceoptions {
          file filename <files number > <size size> <no-stamp | world-readable |
            no-world-readable>;
          flag flag;
        }
      }
    }
  }
```

Related Documentation

- 802.1X for J-EX Series Switches Overview on page 2253
- Example: Configure Automatic VLAN Administration Using GVRP on page 1087
- Understanding MAC RADIUS Authentication on J-EX Series Switches
- Understanding Server Fail Fallback and 802.1X Authentication on J-EX Series Switches on page 2258
- IGMP Snooping on J-EX Series Switches Overview on page 2047
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261
- Understanding MSTP for J-EX Series Switches on page 1277
- Understanding Multiple VLAN Registration Protocol (MVRP) on J-EX Series Switches on page 1054
- Understanding Ethernet OAM Connectivity Fault Management for a J-EX Series Switch on page 3463
- Understanding Ethernet OAM Link Fault Management for a J-EX Series Switch on page 3427
- Understanding RSTP for J-EX Series Switches on page 1276
- Understanding STP for J-EX Series Switches on page 1275
- Understanding How to Use sFlow Technology for Network Monitoring on a J-EX Series Switch on page 3283
- Understanding VSTP for J-EX Series Switches on page 1281

collector

Syntax

```
collector {
  ip-address;
  udp-port port-number;
}
```

Hierarchy Level [edit protocols **sflow**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a remote collector for sFlow network traffic monitoring. The switch sends sFlow UDP datagrams to this collector for analysis. You can configure up to four collectors on the switch. You configure a collector by specifying its IP address and a UDP port.

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [edit protocols] Configuration Statement Hierarchy on page 48
- Example: Monitoring Network Traffic Using sFlow Technology on J-EX Series Switches on page 3285
- Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 3290

disable

Syntax

```
disable;
```

Hierarchy Level [edit protocols **sflow**],
[edit protocols **sflow interfaces** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable the sFlow monitoring protocol on all interfaces on the switch or on the specified interface.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [edit protocols] Configuration Statement Hierarchy on page 48
- Example: Monitoring Network Traffic Using sFlow Technology on J-EX Series Switches on page 3285
- Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 3290

interfaces

Syntax

```
interfaces interface-name {
    disable;
    polling-interval seconds;
    sample-rate number;
}
```

Hierarchy Level [edit protocols sflow]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure sFlow network traffic monitoring on the specified interface on the switch. You can configure sFlow parameters such as polling interval and sample rate with different values on different interfaces, and you can also disable sFlow monitoring on individual interfaces.

The remaining statements are explained separately.

Options ***interface-name***—Name of the interface on which to configure sFlow parameters.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [edit protocols] Configuration Statement Hierarchy on page 48
- Example: Monitoring Network Traffic Using sFlow Technology on J-EX Series Switches on page 3285
- Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 3290

polling-interval

Syntax polling-interval *seconds*;

Hierarchy Level [edit protocols **sflow**],
[edit protocols **sflow interfaces** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the interval (in seconds) that the switch waits between port statistics update messages. "Polling" refers to the switch's gathering various statistics for the network interfaces configured for sFlow monitoring and exporting the statistics to the configured sFlow collector.

Default If no polling interval is configured for a particular interface, the switch waits the number of seconds that is configured for the global sFlow configuration. If no global interval is configured, the switch waits 20 seconds between messages.

Options ***seconds***—Number of seconds between port statistics update messages. A **0** (zero) value specifies that polling is disabled.

Range: 0–3600 seconds

Default: 20 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [edit protocols] Configuration Statement Hierarchy on page 48
- Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 3290
- Example: Monitoring Network Traffic Using sFlow Technology on J-EX Series Switches on page 3285

sample-rate

Syntax sample-rate *number*;

Hierarchy Level [edit protocols **sflow**],
[edit protocols **sflow interfaces** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set the ratio of the number of packets to be sampled in sFlow network traffic monitoring. For example, if you specify a rate of 1000, every thousandth packet (1 packet out of 1000) is sampled.

Default If no sample rate is configured for a particular interface, the switch samples at the rate configured for the global sFlow configuration. If no global rate is configured, the switch samples 1 in 2000 packets.

Options *number*—Denominator of the ratio that composes the sample rate.

Range: 100–1,048,576

Default: 2000

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [edit protocols] Configuration Statement Hierarchy on page 48
- Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 3290
- Example: Monitoring Network Traffic Using sFlow Technology on J-EX Series Switches on page 3285

sflow

Syntax

```
sflow {
  agent-id ip-address;
  collector {
    ip-address;
    udp-port port-number;
  }
  disable;
  interfaces interface-name {
    disable;
    polling-interval seconds;
    sample-rate number;
  }
  polling-interval seconds;
  sample-rate number;
  source-ip ip-address;
}
```

Hierarchy Level [edit protocols]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.
Options **agent-id** and **source-ip** added in Junos OS Release 10.2 for J-EX Series switches.

Description Configure sFlow technology, designed for monitoring high-speed switched or routed networks, to continuously monitor traffic at wire speed on specified interfaces simultaneously. sFlow data can be used to provide network traffic visibility information.

The remaining statements are explained separately.

Default The sFlow protocol is disabled by default.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [edit protocols] Configuration Statement Hierarchy on page 48
- Example: Monitoring Network Traffic Using sFlow Technology on J-EX Series Switches on page 3285
- Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 3290

udp-port

Syntax udp-port *port-number*;

Hierarchy Level [edit protocols **sflow collector**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the UDP port for a remote collector for sFlow network traffic monitoring. The switch sends sFlow UDP datagrams to the collector for analysis.

Options ***port-number***—UDP port number for this collector.
Default: 6343

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- [edit protocols] Configuration Statement Hierarchy on page 48
- Example: Monitoring Network Traffic Using sFlow Technology on J-EX Series Switches on page 3285
- Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 3290

Operational Mode Commands for sFlow Technology

show sflow

Syntax

show sflow
<collector>
<interface>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display default sFlow technology configuration information.

Options

none—Display default sFlow technology configuration information.

collector—(Optional) Display standard status information about the specified sFlow collector.

interface—(Optional) Display standard status information about the specified sFlow interface.

Required Privilege Level

view

Related Documentation

- **show sflow interface on page 3308**
- **show sflow collector on page 3307**
- Example: Monitoring Network Traffic Using sFlow Technology on J-EX Series Switches on page 3285
- Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 3290

List of Sample Output

show sflow on page 3305

Output Fields

Table 440 on page 3305 lists the output fields for the **show sflow** command. Output fields are listed in the approximate order in which they appear.

Table 440: show sflow Output Fields

Field Name	Field Description	Level of Output
sFlow	Status of the feature: **enabled** or **disabled**.	All levels
Sample rate	Rate at which packets are sampled.	All levels
Sample limit	Number of packets sampled per second. The sample limit cannot be configured and is set to 300 packets/second.	All levels
Polling interval	Interval at which the sFlow agent polls the interface.	All levels
Agent ID	The IP address assigned to the sFlow agent.	All levels

show sflow

```
sFlow         : Enabled
Sample rate      : 1:1000
```

```
Sample limit     : 300 packets/second
Polling interval : 20 seconds
Agent ID    : 10.93.54.7
```

show sflow collector

Syntax show sflow collector

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Displays a list of configured sFlow collectors and their properties.

Required Privilege Level view

Related Documentation

- **show sflow on page 3305**
- **show sflow interfaces on page 3308**
- Example: Monitoring Network Traffic Using sFlow Technology on J-EX Series Switches on page 3285
- Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 3290

Output Fields Table 441 on page 3307 lists the output fields for the **show sflow collector** command. Output fields are listed in the approximate order in which they appear.

Table 441: show sflow collector Output Fields

Field Name	Field Description	Level of Output
IP address	IP address of the collector.	All levels
UDP port	UDP port number.	All levels
No of samples	Packet sampling rate.	All levels

show sflow collector

```
IP-address     UDP-Port  No of samples
10.204.32.46    5600    1000
100.204.32.76    3400    1000
```

show sflow interface

Syntax show sflow interface

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display the interfaces on which sFlow technology is enabled and the sampling parameters.

Required Privilege Level view

Related Documentation

- **show sflow on page 3305**
- **show sflow collector on page 3307**
- Example: Monitoring Network Traffic Using sFlow Technology on J-EX Series Switches on page 3285
- Configuring sFlow Technology for Network Monitoring (CLI Procedure) on page 3290

List of Sample Output **show sflow interface on page 3308**

Output Fields Table 442 on page 3308 lists the output fields for the **show sflow interface** command. Output fields are listed in the approximate order in which they appear.

Table 442: show sflow interface Output Fields

Field Name	Field Description	Level of Output
Interfaces	Interfaces on which sFlow technology is enabled.	All levels
Sample-rate	Rate at which packets are sampled.	All levels
Sample-limit	Number of packets sampled per second. Sample limit cannot be configured and is set to 300 packets/second.	All levels
Polling-interval	The interval at which the sFlow agent polls the interface.	All levels

show sflow interface

```
Interfaces    Sample-rate    Adapted sample rate  Polling-interval
ge-0/0/0.0     1:1000        16000                20 sec
```

CHAPTER 128

SNMP

Configuring SNMP

Configuring SNMP (J-Web Procedure)

You can use the J-Web interface to define system identification information, create SNMP communities, create SNMP trap groups, and configure health monitor options for J-EX Series switches.

To configure SNMP features:

1. Select **Configure** > **Services** > **SNMP**.
2. Enter information into the configuration page for SNMP as described in Table 443 on page 3309.
3. To apply the configuration click **Apply**.

NOTE: After you make changes to the configuration in this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options** > **Commit**. See "Using the Commit Options to Commit Configuration Changes (J-Web Procedure)" on page 334 for details about all commit options.

Table 443: SNMP Configuration Page

Field	Function	Your Action
Identification		
Contact Information	Free-form text string that specifies an administrative contact for the system.	Type contact information for the administrator of the system (such as name and phone number).

Table 443: SNMP Configuration Page *(continued)*

Field	Function	Your Action
System Description	Free-form text string that specifies a description for the system.	Type information that describes the system
Local Engine ID	Provides an administratively unique identifier of an SNMPv3 engine for system identification. The local engine ID contains a prefix and a suffix. The prefix is formatted according to specifications defined in RFC 3411. The suffix is defined by the local engine ID. Generally, the local engine ID suffix is the MAC address of Ethernet management port 0.	Type the MAC address of Ethernet management port 0.
System Location	Free-form text string that specifies the location of the system.	Type location information for the system (lab name or rack name, for example).
System Override Name	Free-form text string that overrides the system hostname.	Type the hostname of the system.
Communities To add a community, click **Add**		
Community Name	Specifies the name of the SNMP community.	Type the name of the community being added.
Authorization	Specifies the type of authorization (either read-only or read-write) for the SNMP community being configured.	Select the authorization (either read-only or read-write) from the list.
Traps To add a trap group, click **Add**.		
Trap Group Name	Specifies the name of the SNMP trap group being configured.	Type the name of the group being added.

Table 443: SNMP Configuration Page *(continued)*

Field	Function	Your Action
Categories	Specifies which trap categories are added to the trap group being configured.	• To generate traps for authentication failures, select **Authentication**. • To generate traps for chassis and environment notifications, select **Chassis**. • To generate traps for configuration changes, select **Configuration**. • To generate traps for link-related notifications (up-down transitions), select **Link**. • To generate traps for remote operation notifications, select **Remote operations**. • To generate traps for remote network monitoring (RMON), select **RMON alarm**. • To generate traps for routing protocol notifications, select **Routing**. • To generate traps on system warm and cold starts, select **Startup**. • To generate traps on Virtual Router Redundancy Protocol (VRRP) events (such as new-master or authentication failures), select **VRRP events**.
Targets	Specifies one or more hostnames or IP addresses for the systems to receive SNMP traps generated by the trap group being configured.	1. Enter the hostname or IP address, in dotted decimal notation, of the target system to receive the SNMP traps. 2. Click **Add**.
Health Monitoring		
Enable Health Monitoring	Enables the SNMP health monitor on the switch. The health monitor periodically (over the time you specify in the interval field) checks the following key indicators of switch health: • Percentage of file storage used • Percentage of Routing Engine CPU used • Percentage of Routing Engine memory used • Percentage of memory used for each system process • Percentage of CPU used by the forwarding process • Percentage of memory used for temporary storage by the forwarding process	Select the check box to enable the health monitor and configure options. Clear the check box to disable the health monitor. NOTE: If you select the **Enable Health Monitoring** check box and do not specify options, then SNMP health monitoring is enabled with default values.
Interval	Specifies the sampling frequency, in seconds, over which the key health indicators are sampled and compared with the rising and falling thresholds. For example, if you configure the interval as 100 seconds, the values are checked every 100 seconds.	Enter an interval time, in seconds, from **1** through **2147483647**. The default value is 300 seconds (5 minutes).

Table 443: SNMP Configuration Page *(continued)*

Field	Function	Your Action
Rising Threshold	Specifies the value at which SNMP generates an event (trap and system log message) when the value of a sampled indicator is increasing. For example, if the rising threshold is 90 (the default), SNMP generates an event when the value of any key indicator reaches or exceeds 90 percent.	Enter a value from **0** through **100**. The default value is **90**.
Falling Threshold	Specifies the value at which SNMP generates an event (trap and system log message) when the value of a sampled indicator is decreasing. For example, if the falling threshold is 80 (the default), SNMP generates an event when the value of any key indicator falls back to 80 percent or less.	Enter a value from **0** through **100**. The default value is **80**. NOTE: The falling threshold value must be less than the rising threshold value.

Related Documentation

- Monitoring System Process Information on page 554
- Monitoring System Properties on page 550

Configuration Statements for SNMP

- [edit snmp] Configuration Statement Hierarchy on page 3312

[edit snmp] Configuration Statement Hierarchy

```
snmp {
    rmon {
        history index {
            bucket-size number;
            interface interface-name;
            interval seconds;
            owner owner-name;
        }
    }
}
```

Related Documentation

- Configuring SNMP (J-Web Procedure) on page 3309
- *Junos OS Network Management Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

address

Syntax address *address*;

Hierarchy Level [edit snmp v3 **target-address** *target-address-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the SNMP target address.

Options ***address***—IPv4 address of the system to receive traps or informs. You must specify an address, not a hostname.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the Address

address-mask

Syntax address-mask *address-mask*;

Hierarchy Level [edit snmp v3 **target-address** *target-address-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Verify the source addresses for a group of target addresses.

Options ***address-mask*** combined with the address defines a range of addresses.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the Address Mask

agent-address

Syntax agent-address outgoing-interface;

Hierarchy Level [edit snmp **trap-options**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set the agent address of all SNMPv1 traps generated by this router. Currently, the only option is **outgoing-interface**, which sets the agent address of each SNMPv1 trap to the address of the outgoing interface of that trap.

Options

outgoing-interface—Value of agent address of all SNMPv1 traps generated by this router. The **outgoing-interface** option sets the agent address of each SNMPv1 trap to the address of the outgoing interface of that trap.

Default: disabled (The agent address is not specified in SNMPv1 traps.)

Required Privilege Level

snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- Configuring the Agent Address for SNMP Traps

alarm

Syntax

```
alarm index {
  description description;
  falling-event-index index;
  falling-threshold integer;
  falling-threshold-interval seconds;
  interval seconds;
  request-type (get-next-request | get-request | walk-request);
  rising-event-index index;
  rising-threshold integer;
  sample-type (absolute-value | delta-value);
  startup-alarm (falling-alarm | rising-alarm | rising-or-falling alarm);
  syslog-subtag syslog-subtag;
  variable oid-variable;
}
```

Hierarchy Level [edit snmp rmon]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure RMON alarm entries.

Options *index*—Identifies this alarm entry as an integer.

The remaining statements are explained separately.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- Configuring an Alarm Entry and Its Attributes
- **event on page 3326**

authorization

Syntax authorization *authorization*;

Hierarchy Level [edit snmp **community** *community-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set the access authorization for SNMP **Get**, **GetBulk**, **GetNext**, and **Set** requests.

Options ***authorization***—Access authorization level:

- **read-only**—Enable **Get**, **GetNext**, and **GetBulk** requests.
- **read-write**—Enable all requests, including **Set** requests. You must configure a view to enable **Set** requests.

Default: read-only

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- Configuring the SNMP Community String

bucket-size

Syntax bucket-size *number*;

Hierarchy Level [edit **snmp rmon history**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the sampling of Ethernet statistics for network fault diagnosis, planning, and performance tuning.

Default 50

Options ***number***—Number of discrete samples of Ethernet statistics requested.

Required Privilege Level snmp—To view this statement in the configuration.
snmp–control—To add this statement to the configuration.

Related Documentation

- Configuring SNMP (J-Web Procedure) on page 3309
- *Junos OS Network Management Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

categories

Syntax

```
categories {
  category;
}
```

Hierarchy Level [edit snmp **trap-group** *group-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define the types of traps that are sent to the targets of the named trap group.

Default If you omit the **categories** statement, all trap types are included in trap notifications.

Options ***category***—Name of a trap type.

Values: **authentication**, **chassis**, **configuration**, **link**, **remote-operations**, **rmon-alarm**, **routing**, **sonet-alarms**, **startup**, **vrrp-events**

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- Configuring SNMP Trap Groups

client-list

Syntax

```
client-list client-list-name {
  ip-addresses;
}
```

Hierarchy Level [edit snmp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define a list of SNMP clients.

Options ***client-list-name***—Name of the client list.

ip-addresses—IP addresses of the SNMP clients to be added to the client list,

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- Adding a Group of Clients to an SNMP Community

client-list-name

Syntax client-list-name *client-list-name*;

Hierarchy Level [edit snmp **community** *community-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Add a client list or prefix list to an SNMP community.

Options ***client-list-name***—Name of the client list or prefix list.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Adding a Group of Clients to an SNMP Community

clients

Syntax
```
clients {
  address <restrict>;
}
```

Hierarchy Level [edit snmp **community** *community-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the IPv4 or IPv6 addresses of the SNMP client hosts that are authorized to use this community.

Default If you omit the **clients** statement, all SNMP clients using this community string are authorized to access the router.

Options ***address***—Address of an SNMP client that is authorized to access this router. You must specify an address, not a hostname. To specify more than one client, include multiple ***address*** options.

restrict—(Optional) Do not allow the specified SNMP client to access the router.

Default: The client is granted access.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the SNMP Community String

commit-delay

Syntax commit-delay *seconds*;

Hierarchy Level [edit snmp **nonvolatile**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the timer for the SNMP **Set** reply and start of the commit.

Options ***seconds***—Delay between affirmative SNMP **Set** reply and start of the commit.
Default: 5 seconds

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the Commit Delay Timer

community

Syntax

```
community community-name {
    authorization authorization;
    client-list-name client-list-name;
    clients {
        address restrict;
    }
    view view-name;
}
```

Hierarchy Level [edit snmp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define an SNMP community. An SNMP community authorizes SNMP clients based on the source IP address of incoming SNMP request packets. A community also defines which MIB objects are available and the operations (read-only or read-write) allowed on those objects.

The SNMP client application specifies an SNMP community name in **Get**, **GetBulk**, **GetNext**, and **Set** SNMP requests.

Default If you omit the **community** statement, all SNMP requests are denied.

Options ***community-name***—Community string. If the name includes spaces, enclose it in quotation marks (" ").

The remaining statements are explained separately.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- Configuring the SNMP Community String

community

Syntax community *community-name*;

Hierarchy Level [edit snmp rmon **event** *index*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description The trap group that is used when generating a trap (if **eventType** is configured to send traps). If that trap group has the **rmon-alarm** trap category configured, a trap is sent to all the targets configured for that trap group. The community string in the trap matches the name of the trap group (and hence, the value of **eventCommunity**). If nothing is configured, traps are sent to each group with the **rmon-alarm** category set.

Options ***community-name***—Identifies the trap group that is used when generating a trap if the event is configured to send traps.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring an Event Entry and Its Attributes

community-name

Syntax community-name *community-name*;

Hierarchy Level [edit snmp v3 **snmp-community** *community-index*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description The community name defines an SNMP community. The SNMP community authorizes SNMPv1 or SNMPv2 clients. The access privileges associated with the configured security name define which MIB objects are available and the operations (notify, read, or write) allowed on those objects.

Options ***community-name***—Community string for an SNMPv1 or SNMPv2c community. If unconfigured, it is the same as the community index. If the name includes spaces, enclose it in quotation marks (" ").

NOTE: Community names must be unique. You cannot configure the same community name at the [edit snmp community] and [edit snmp v3 snmp-community *community-index*] hierarchy levels.

The community name at the [edit snmp v3 snmp-community *community-index*] hierarchy level is encrypted and not displayed in the command-line interface (CLI).

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- Configuring the SNMPv3 Community

contact

Syntax contact *contact*;

Hierarchy Level [edit snmp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define the value of the MIB II **sysContact** object, which is the contact person for the managed system.

Options ***contact***—Name of contact person. If the name includes spaces, enclose it in quotation marks (" ").

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the System Contact on a Device Running Junos OS

description

Syntax description *description*;

Hierarchy Level [edit snmp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define the value of the MIB II **sysDescription** object, which is the description of the system being managed.

Options ***description***—System description. If the name includes spaces, enclose it in quotation marks (" ").

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the System Description on a Device Running Junos OS

description

Syntax description *description*;

Hierarchy Level [edit snmp rmon **alarm** *index*],
[edit snmp rmon **event** *index*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Text description of alarm or event.

Options ***description***—Text description of an alarm or event entry. If the description includes spaces, enclose it in quotation marks (" ").

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the Description
- Configuring an Event Entry and Its Attributes

destination-port

Syntax destination-port *port-number*;

Hierarchy Level [edit snmp **trap-group**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Assign a trap port number other than the default.

Default If you omit this statement, the default port is 162.

Options ***port-number***—SNMP trap port number.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring SNMP Trap Groups

engine-id

Syntax

```
engine-id {
   (local engine-id-suffix | use-default-ip-address | use-mac-address);
}
```

Hierarchy Level [edit snmp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description The local engine ID is defined as the administratively unique identifier of an SNMPv3 engine, and is used for identification, not for addressing. There are two parts of an engine ID: prefix and suffix. The prefix is formatted according to the specifications defined in RFC 3411, *An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks*. You can configure the suffix here.

NOTE: SNMPv3 authentication and encryption keys are generated based on the associated passwords and the engine ID. If you configure or change the engine ID, you must commit the new engine ID before you configure SNMPv3 users. Otherwise the keys generated from the configured passwords are based on the previous engine ID.

For the engine ID, we recommend using the MAC address of fxp0.

Options

local *engine-id-suffix*—Explicit setting for the engine ID suffix.

use-default-ip-address—The engine ID suffix is generated from the default IP address.

use-mac-address—The SNMP engine identifier is generated from the MAC address of the management interface on the router.

Default: use-default-ip-address

Required Privilege Level

snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- Configuring the Local Engine ID

event

Syntax

```
event index {
    community community-name;
    description description;
    type type;
}
```

Hierarchy Level [edit snmp rmon]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure RMON event entries.

Options *index*—Identifier for a specific event entry.

The remaining statements are explained separately.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- Configuring an Event Entry and Its Attributes
- **alarm on page 3315**

falling-event-index

Syntax falling-event-index *index*;

Hierarchy Level [edit snmp rmon alarm *index*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description The index of the event entry that is used when a falling threshold is crossed. If this value is zero, no event is triggered.

Options *index*—Index of the event entry that is used when a falling threshold is crossed.

Range: 0 through 65,535

Default: 0

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- Configuring the Falling Event Index or Rising Event Index
- **rising-event-index on page 3343**

falling-threshold

Syntax falling-threshold *percentage*;

Hierarchy Level [edit snmp **health-monitor**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description The lower threshold is expressed as a percentage of the maximum possible value for the sampled variable. When the current sampled value is less than or equal to this threshold, and the value at the last sampling interval is greater than this threshold, a single event is generated. A single event is also generated if the first sample after this entry becomes valid is less than or equal to this threshold. After a falling event is generated, another falling event cannot be generated until the sampled value rises above this threshold and reaches the **rising-threshold**.

Options *percentage*—The lower threshold for the alarm entry.

Range: 1 through 100

Default: 70 percent of the maximum possible value

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- Configuring the Falling Threshold or Rising Threshold
- **rising-threshold on page 3344**

falling-threshold

Syntax falling-threshold *integer*;

Hierarchy Level [edit snmp rmon **alarm** *index*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description The lower threshold for the sampled variable. When the current sampled value is less than or equal to this threshold, and the value at the last sampling interval is greater than this threshold, a single event is generated. A single event is also generated if the first sample after this entry becomes valid is less than or equal to this threshold, and the associated startup-alarm value is equal to falling-alarm value or rising-or-falling-alarm value. After a falling event is generated, another falling event cannot be generated until the sampled value rises above this threshold and reaches the **rising-threshold**.

Options *integer*—The lower threshold for the alarm entry.
Range: -2,147,483,648 through 2,147,483,647
Default: 20 percent less than **rising-threshold**

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the Falling Threshold or Rising Threshold
- **rising-threshold on page 3345**

falling-threshold-interval

Syntax falling-threshold-interval *seconds*;

Hierarchy Level [edit snmp rmon **alarm** *index*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Interval between samples when the rising threshold is crossed. Once the alarm crosses the falling threshold, the regular sampling interval is used.

Options *seconds*—Time between samples, in seconds.
Range: 1 through 2,147,483,647 seconds
Default: 60 seconds

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the Falling Threshold Interval
- **interval on page 3335**

filter-duplicates

Syntax filter-duplicates;

Hierarchy Level [edit snmp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Filter duplicate **Get**, **GetNext**, or **GetBulk** SNMP requests.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Filtering Duplicate SNMP Requests

filter-interfaces

Syntax

```
filter-interfaces {
  interfaces {
    all-internal-interfaces;
    interface 1;
    interface 2;
  }
}
```

Hierarchy Level [edit snmp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Filter out information related to specific interfaces from the output of SNMP **Get** and **GetNext** requests performed on interface-related MIBs.

Options

all-internal-interfaces—Filters out information related to internal interfaces from the output of SNMP **Get** and **GetNext** requests.

interfaces—Specifies the interfaces to filter out from the output of SNMP **Get** and **GetNext** requests.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Filtering Interface Information Out of SNMP Get and GetNext Output

group (Configuring Group Name)

Syntax group *group-name*;

Hierarchy Level [edit snmp v3 vacm access]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Assign the security name to a group.

Options ***group-name***—SNMPv3 group name created for the SNMPv3 group.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the Group

group (Defining Access Privileges for an SNMPv3 Group)

Syntax group *group-name*;

Hierarchy Level [edit snmp v3 vacm security-to-group security-model (usm | v1 | v2c) security-name *security-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define access privileges granted to a group.

Options ***group-name***—Identifies a collection of SNMP security names that belong to the same access policy SNMP.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the Group

health-monitor

Syntax

```
health-monitor {
    falling-threshold percentage;
    interval seconds;
    rising-threshold percentage;
}
```

Hierarchy Level [edit snmp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure health monitoring.

The remaining statements are explained separately.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- Configuring Health Monitoring on Devices Running Junos OS

history

Syntax

```
history history-index {
    bucket-size number;
    interface interface-name;
    interval seconds;
    owner owner-name;
}
```

Hierarchy Level [edit snmp rmon]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure RMON history group entries. This RMON feature can be used with the Simple Network Management Protocol (SNMP) agent in the switch to monitor all the traffic flowing among switches on all connected LAN segments. It collects statistics in accordance with user-configurable parameters.

The history group controls the periodic statistical sampling of data from various types of networks. This group contains configuration entries that specify an interface, polling period, and other parameters. The **interface *interface-name*** statement is mandatory. Other statements in the history group are optional.

Default Not configured.

Options ***history-index***—Identifies this history entry as an integer.

Range: 1 through 655535

Required Privilege Level snmp—To view this statement in the configuration.
snmp–control—To add this statement to the configuration.

Related Documentation

- Configuring SNMP (J-Web Procedure) on page 3309
- *Junos OS Network Management Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

interface

Syntax interface [*interface-names*];

Hierarchy Level [edit snmp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the interfaces on which SNMP requests can be accepted.

Default If you omit this statement, SNMP requests entering the router through any interface are accepted.

Options ***interface-names***—Names of one or more logical interfaces.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the Interfaces on Which SNMP Requests Can Be Accepted

interface

Syntax interface *interface-name*;

Hierarchy Level [edit **snmp rmon history** *history-index*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the interface to be monitored in the specified RMON history entry.

Only one interface can be specified for a particular RMON history index. There is a one-to-one relationship between the interface and the history index. The interface must be specified in order for the RMON history to be created.

Options ***interface-name***—Specify the interface to be monitored within the specified entry of the RMON history of Ethernet statistics.

Required Privilege Level snmp—To view this statement in the configuration.
snmp–control—To add this statement to the configuration.

Related Documentation
- Configuring SNMP (J-Web Procedure) on page 3309
- *Junos OS Network Management Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

interval

Syntax interval *seconds*;

Hierarchy Level [edit **snmp rmon history**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the interval over which data is to be sampled for the specified interface.

Default **1800 sec**

Options ***seconds***—Interval at which data is to be sampled for the specified interface.

Required Privilege Level snmp—To view this statement in the configuration.
snmp–control—To add this statement to the configuration.

interval

Syntax interval *seconds*;

Hierarchy Level [edit snmp **health-monitor**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Interval between samples.

Options ***seconds***—Time between samples, in seconds.
Range: **1** through **2147483647** seconds
Default: **300** seconds

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- Configuring the Interval

interval

Syntax interval *seconds*;

Hierarchy Level [edit snmp rmon **alarm** *index*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Interval between samples.

Options ***seconds***—Time between samples, in seconds.
Range: 1 through 2,147,483,647 seconds
Default: 60 seconds

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the Interval

location

Syntax location *location*;

Hierarchy Level [edit snmp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define the value of the MIB II **sysLocation** object, which is the physical location of the managed system.

Options ***location***—Location of the local system. You must enclose the name within quotation marks (" ").

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the System Location for a Device Running Junos OS

logical-system

Syntax
```
logical-system logical-system-name {
  routing-instance routing-instance-name;
}
```

Hierarchy Level
[edit snmp **community** *community-name*],
[edit snmp **trap-group**],
[edit snmp **trap-options**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify a logical system name for SNMP v1 and v2c clients.

Options
logical-system-name—Name of the logical system.

routing-instance ***routing-instance-name***—Statement to specify a routing instance associated with the logical system.

Required Privilege Level
snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Specifying a Routing Instance in an SNMPv1 or SNMPv2c Community

message-processing-model

Syntax
```
message-processing-model (v1 | v2c | v3);
```

Hierarchy Level
[edit snmp v3 target-parameters *target-parameter-name* **parameters**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the message processing model to be used when generating SNMP notifications.

Options
v1—SNMPv1 message process model.

v2c—SNMPv2c message process model.

v3—SNMPv3 message process model.

Required Privilege Level
snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the Message Processing Model

name

Syntax name *name*;

Hierarchy Level [edit snmp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set the system name from the command-line interface.

Options *name*—System name override.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the System Name

nonvolatile

Syntax
```
nonvolatile {
  commit-delay seconds;
}
```

Hierarchy Level [edit snmp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure options for SNMP **Set** requests.

The statement is explained separately.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the Commit Delay Timer
- **commit-delay on page 3319**

notify

Syntax

```
notify name {
  tag tag-name;
  type (trap | inform);
}
```

Hierarchy Level [edit snmp v3]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Select management targets for notifications as well as the type of notifications. Notifications can be either traps or informs.

Options

name—Name assigned to the notification.

tag-name—Notifications are sent to all targets configured with this tag.

type—Notification type is **trap** or **inform**. Traps are unconfirmed notifications. Informs are confirmed notifications.

Required Privilege Level
snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- Configuring the Inform Notification Type and Target Address
- Configuring the SNMPv3 Trap Notification

notify-filter (Configuring the Profile Name)

Syntax

```
notify-filter profile-name {
  oid oid (include | exclude);
}
```

Hierarchy Level [edit snmp v3]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define a group of MIB objects on which to define access. The notify filter limits the type of traps or informs sent to the NMS.

Options

profile-name—Name assigned to the notify filter.

The remaining statement is explained separately.

Required Privilege Level
snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- Configuring the Trap Notification Filter
- **oid on page 3340**

notify-filter (Applying to the Management Target)

Syntax notify-filter *profile-name*;

Hierarchy Level [edit snmp v3 **target-parameters** *target-parameters-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the notify filter to be used by a specific set of target parameters.

Options ***profile-name***—Name of the notify filter to apply to notifications.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Applying the Trap Notification Filter

notify-view

Syntax notify-view *view-name*;

Hierarchy Level [edit snmp v3 vacm access group *group-name* default-context-prefix security-model (any | usm | v1 | v2c) **security-level** (authentication | none | privacy)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Associate the view with a community or a group name (SNMPv3).

Options ***view-name***—Name of the view to which the SNMP user group has access.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring MIB Views
- Configuring the Notify View

oid

Syntax oid *object-identifier* (exclude| include);

Hierarchy Level [edit snmp view *view-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify an object identifier (OID) used to represent a subtree of MIB objects.

Options

exclude—Exclude the subtree of MIB objects represented by the specified OID.

include—Include the subtree of MIB objects represented by the specified OID.

object-identifier—OID used to represent a subtree of MIB objects. All MIB objects represented by this statement have the specified OID as a prefix. You can specify the OID using either a sequence of dotted integers or a subtree name.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- Configuring MIB Views

oid

Syntax oid *oid* (include | exclude);

Hierarchy Level [edit snmp v3 **notify-filter** *profile-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify an object identifier (OID) used to represent a subtree of MIB objects.

Options

exclude—Exclude the subtree of MIB objects represented by the specified OID.

include—Include the subtree of MIB objects represented by the specified OID.

oid—Object identifier used to represent a subtree of MIB objects. All MIB objects represented by this statement have the specified OID as a prefix. You can specify the OID using either a sequence of dotted integers or a subtree name.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- Configuring the Trap Notification Filter

owner

Syntax owner *owner-name*;

Hierarchy Level [edit **snmp rmon history**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the user or group responsible for this configuration.

Options ***owner-name***—The user or group responsible for this configuration.
Range: 0 through 32 alphanumeric characters

Required Privilege Level snmp—To view this statement in the configuration.
snmp–control—To add this statement to the configuration.

Related Documentation
- Configuring SNMP (J-Web Procedure) on page 3309
- *Junos OS Network Management Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

parameters

Syntax
```
parameters {
  message-processing-model (v1 | v2c | v3);
  security-level (none | authentication | privacy);
  security-model (usm | v1 | v2c);
  security-name security-name;
}
```

Hierarchy Level [edit snmp v3 **target-parameters** *target-parameters-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a set of target parameters.

The remaining statements are explained separately.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Defining and Configuring the Trap Target Parameters

port

Syntax port *port-number*;

Hierarchy Level [edit snmp v3 **target-address** *target-address-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a UDP port number for an SNMP target.

Default If you omit this statement, the default port is 162.

Options ***port-number***—Port number for the SNMP target.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the Port

read-view

Syntax read-view *view-name*;

Hierarchy Level [edit snmp v3 vacm access group *group-name* default-context-prefix security-model (any | usm | v1 | v2c) **security-level** (authentication | none | privacy)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Associate the view with a community or a group name (SNMPv3).

Options ***view-name***—The name of the view to which the SNMP user group has access.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the Read View
- Configuring MIB Views

request-type

Syntax request-type (get-next-request | get-request | walk-request);

Hierarchy Level [edit snmp rmon **alarm** *index*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Extend monitoring to a specific SNMP object instance (**get-request**), or extend monitoring to all object instances belonging to a MIB branch (**walk-request**), or extend monitoring to the next object instance after the instance specified in the configuration (**get-next-request**).

Options **get-next-request**—Performs an SNMP get next request.

get-request—Performs an SNMP get request.

walk-request—Performs an SNMP walk request.

Default: **walk-request**

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the Request Type
- **variable on page 3368**

rising-event-index

Syntax rising-event-index *index*;

Hierarchy Level [edit snmp rmon **alarm** *index*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Index of the event entry that is used when a rising threshold is crossed. If this value is zero, no event is triggered.

Options *index*—Index of the event entry that is used when a rising threshold is crossed.

Range: 0 through 65,535

Default: 0

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the Falling Event Index or Rising Event Index
- **falling-event-index on page 3326**

rising-threshold

Syntax rising-threshold *percentage*;

Hierarchy Level [edit snmp **health-monitor**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description The upper threshold is expressed as a percentage of the maximum possible value for the sampled variable. When the current sampled value is greater than or equal to this threshold, and the value at the last sampling interval is less than this threshold, a single event is generated. A single event is also generated if the first sample after this entry becomes valid is greater than or equal to this threshold. After a rising event is generated, another rising event cannot be generated until the sampled value falls below this threshold and reaches the **falling-threshold**.

Options *percentage*—The lower threshold for the alarm entry.

Range: 1 through 100

Default: 80 percent of the maximum possible value

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- **falling-threshold on page 3327**
- Configuring the Falling Threshold or Rising Threshold

rising-threshold

Syntax rising-threshold *integer*;

Hierarchy Level [edit snmp rmon **alarm** *index*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Upper threshold for the sampled variable. When the current sampled value is greater than or equal to this threshold, and the value at the last sampling interval is less than this threshold, a single event is generated. A single event is also generated if the first sample after this entry becomes valid is greater than or equal to this threshold, and the associated startup alarm value is equal to the falling alarm or rising or falling alarm value. After a rising event is generated, another rising event cannot be generated until the sampled value falls below this threshold and reaches the falling threshold.

Options ***integer***—The lower threshold for the alarm entry.
Range: –2,147,483,648 through 2,147,483,647

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the Falling Threshold or Rising Threshold
- **falling-threshold on page 3328**

rmon

Syntax rmon { ... }

Hierarchy Level [edit snmp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure Remote Monitoring.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring an Alarm Entry and Its Attributes

rmon

Syntax

```
rmon {
  history history-index {
    interface interface-name;
    bucket-size number;
    interval seconds;
    owner owner-name;
  }
}
```

Hierarchy Level [edit snmp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description RMON is an existing feature of the Junos OS.

The RMON specification provides network administrators with comprehensive network fault diagnosis, planning, and performance tuning information. It delivers this information in nine groups of monitoring elements, each providing specific sets of data to meet common network monitoring requirements. Each group is optional, so that vendors do not need to support all the groups within the MIB.

Junos OS supports RMON Statistics, History, Alarm, and Event groups. The J-EX Series documentation describes only the **rmon history** statement.

The statements are explained separately.

Default Disabled.

Required Privilege Level snmp—To view this statement in the configuration.
snmp–control—To add this statement to the configuration.

Related Documentation

- Configuring SNMP (J-Web Procedure) on page 3309
- *Junos OS Network Management Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

routing-instance

Syntax routing-instance *routing-instance-name*;

Hierarchy Level [edit snmp **community** *community-name*],
[edit snmp **community** *community-name* logical-system *logical-system-name*],
[edit snmp **trap-group** *group*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify a routing instance for SNMPv1 and SNMPv2 trap targets. All targets configured in the trap group use this routing instance.

If the routing instance is defined within a logical system, include the **logical-system** ***logical-system-name*** statement at the **[edit snmp community** ***community-name*****]** hierarchy level and specify the **routing-instance** statement under the **[edit snmp community** ***community-name*** **logical-system** ***logical system-name*****]** hierarchy level.

Options ***routing-instance-name***—Name of the routing instance.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- Configuring SNMP Trap Groups
- Configuring the Source Address for SNMP Traps
- Specifying a Routing Instance in an SNMPv1 or SNMPv2c Community

routing-instance

Syntax routing-instance *routing-instance-name*;

Hierarchy Level [edit snmp v3 **target-address** *target-address-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify a routing instance for an SNMPv3 trap target.

Options ***routing-instance-name***—Name of the routing instance.

To configure a routing instance within a logical system, specify the logical system name followed by the routing instance name. Use a slash (**/**) to separate the two names (for example, **test-ls/test-ri**). To configure the default routing instance on a logical system, specify the logical system name followed by **default** (for example, **test-ls/default**).

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the Trap Target Address

sample-type

Syntax sample-type (absolute-value | delta-value);

Hierarchy Level [edit snmp rmon **alarm** *index*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Method of sampling the selected variable.

Options **absolute-value**—Actual value of the selected variable is used when comparing against the thresholds.

delta-value—Difference between samples of the selected variable is used when comparing against the thresholds.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the Sample Type

security-level (Generating SNMP Notifications)

Syntax security-level (authentication | none | privacy);

Hierarchy Level [edit snmp v3 target-parameters *target-parameters-name* parameters]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the security level to use when generating SNMP notifications.

Options **authentication**—Provides authentication but no encryption.

none—No authentication and no encryption.

privacy—Provides authentication and encryption.

Default: none

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the Security Level

security-level (Defining Access Privileges)

Syntax
```
security-level (authentication | none | privacy) {
  notify-view view-name;
  read-view view-name;
  write-view view-name;
}
```

Hierarchy Level [edit snmp v3 vacm access **group** *group-name* default-context-prefix **security-model** (any | usm | v1 | v2c)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define the security level used for access privileges.

Options **authentication**—Provides authentication but no encryption.

none—No authentication and no encryption.

privacy—Provides authentication and encryption.

Default: none

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the Security Level

security-model (Access Privileges)

Syntax security-model (usm | v1 | v2c);

Hierarchy Level [edit snmp v3 vacm access group *group-name* default-context-prefix]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a group's security model used for access privileges.

Options **usm**—SNMPv3 security model.

v1—SNMPv1 security model.

v2c—SNMPv2c security model.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the Security Model

security-model (Group)

Syntax
```
security-model (usm | v1 | v2c) {
  security-name security-name {
    group group-name;
  }
}
```

Hierarchy Level [edit snmp v3 vacm **security-to-group**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define a security model for a group.

Options **usm**—SNMPv3 security model.

v1—SNMPv1 security model.

v2c—SNMPv2c security model.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the Security Model

security-model (SNMP Notifications)

Syntax security-model (usm | v1 | v2c);

Hierarchy Level [edit snmp v3 target-parameters *target-parameters-name* parameters]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a group's security model used with sending notifications.

Options **usm**—SNMPv3 security model.

v1—SNMPv1 security model.

v2c—SNMPv2c security model.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the Security Model

security-name (Security Group)

Syntax
```
security-name security-name {
  group group-name;
}
```

Hierarchy Level [edit snmp v3 vacm security-to-group **security-model** (usm | v1 | v2c)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Associate a group or a community string with a configured security group.

Options ***security-name***—Username configured at the **[edit snmp v3 usm local-engine user *username*]** hierarchy level. For SNMPv1 and SNMPv2c, the security name is the community string configured at the **[edit snmp v3 snmp-community *community-index*]** hierarchy level.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Assigning Security Names to Groups

security-name (Community String)

Syntax security-name *security-name*;

Hierarchy Level [edit snmp v3 **snmp-community** *community-index*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Associate the community string configured at the **[edit snmp v3 snmp-community *community-index*]** hierarchy level to a security name.

Options ***security-name***—Name used when performing access control.

NOTE: The security name must match the configured security name at the [edit snmp v3 target-parameters *target-parameters-name* parameters] hierarchy level when you configure traps or informs.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- Configuring the Security Names

security-name (SNMP Notifications)

Syntax security-name *security-name*;

Hierarchy Level [edit snmp v3 target-parameters *target-parameters-name* parameters]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the security name used when generating SNMP notifications.

Options ***security-name***—Identifies the user that is used when generating the notification if the USM security model is used. Identifies the SNMP community used when generating the notification if the v1 or v2c security models are used.

NOTE: The access privileges for the group associated with this security name must allow this notification to be sent.

If you are using the v1 or v2 security models, the security name at the [edit snmp v3 vacm security-to-group] hierarchy level must match the security name at the [edit snmp v3 snmp-community *community-index*] hierarchy level.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the Security Name

security-to-group

Syntax
```
security-to-group {
  security-model (usm | v1 | v2c) {
    group group-name;
    security-name security-name;
  }
}
```

Hierarchy Level [edit snmp v3 vacm]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the group to which a specific security name belongs.

The remaining statements are explained separately.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Assigning Security Names to Groups

snmp

Syntax snmp { ... }

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure SNMP.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring SNMP on a Device Running Junos OS

snmp

Syntax

```
snmp {
  rmon {
    history index {
      interface interface-name;
      bucket-size number;
      interval seconds;
      owner owner-name;
    }
  }
}
```

Hierarchy Level [edit]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure SNMP.

The statements are explained separately.

Required Privilege Level snmp—To view this statement in the configuration.
snmp–control—To add this statement to the configuration.

Related Documentation
- *Junos OS Network Management Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

snmp-community

Syntax

```
snmp-community community-index {
    community-name community-name;
    security-name security-name;
    tag tag-name;
}
```

Hierarchy Level [edit snmp v3]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the SNMP community.

Options ***community-index***—(Optional) String that identifies an SNMP community.

The remaining statements are explained separately.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- Configuring the SNMPv3 Community

source-address

Syntax source-address *address*;

Hierarchy Level [edit snmp trap-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set the source address of every SNMP trap packet sent by this router to a single address regardless of the outgoing interface. If the source address is not specified, the default is to use the address of the outgoing interface as the source address.

Options ***address***—Source address of SNMP traps. You can configure the source address of trap packets two ways: **lo0** or a valid IPv4 address configured on one of the router interfaces. The value **lo0** indicates that the source address of all SNMP trap packets is set to the lowest loopback address configured at interface **lo0**.

Default: disabled (The source address is the address of the outgoing interface.)

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- Configuring the Source Address for SNMP Traps

startup-alarm

Syntax startup-alarm (falling-alarm | rising-alarm | rising-or-falling-alarm);

Hierarchy Level [edit snmp rmon alarm *index*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description The alarm that can be sent upon entry startup.

Options

falling-alarm—Generated if the first sample after the alarm entry becomes active is less than or equal to the falling threshold.

rising-alarm—Generated if the first sample after the alarm entry becomes active is greater than or equal to the rising threshold.

rising-or-falling-alarm—Generated if the first sample after the alarm entry becomes active satisfies either of the corresponding thresholds.

Default: rising-or-falling-alarm

Required Privilege Level
snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the Sample Type

syslog-subtag

Syntax syslog-subtag *syslog-subtag*;

Hierarchy Level [edit snmp rmon alarm *index*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Add a tag to the system log message.

Options

syslog-subtag *syslog-subtag*—Tag of not more than 80 uppercase characters to be added to syslog messages.

Default: None

Required Privilege Level
snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the System Log Tag

tag

Syntax tag *tag-name*;

Hierarchy Level [edit snmp v3 **notify** *name*],
[edit snmp v3 **snmp-community** *community-index*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a set of targets to receive traps or informs (for IPv4 packets only).

Options ***tag-name***—Identifies the address of managers that are allowed to use a community string.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the Tag
- Configuring the SNMPv3 Trap Notification

tag-list

Syntax tag-list *tag-list*;

Hierarchy Level [edit snmp v3 **target-address** *target-address-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure an SNMP tag list used to select target addresses.

Options ***tag-list***—Defines sets of target addresses. To specify more than one tag, specify the tag names as a space-separated list enclosed within double quotes.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the Trap Target Address

target-address

Syntax

```
target-address target-address-name {
    address address;
    address-mask address-mask;
    inform-retry-count number;
    inform-timeout seconds;
    port port-number;
    routing-instance instance;
    tag-list tag-list;
    target-parameters target-parameters-name;
}
```

Hierarchy Level [edit snmp v3]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure a management application's address and parameters to be used in sending notifications.

Options ***target-address-name***—String that identifies the target address.

The remaining statements are explained separately.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- Configuring the Trap Target Address

target-parameters

Syntax

```
target-parameters target-parameters-name {
  profile-name;
  parameters {
    message-processing-model (v1 | v2c | V3);
    security-level (authentication | none | privacy);
    security-model (usm | v1 | v2c);
    security-name security-name;
  }
}
```

Hierarchy Level [edit snmp v3]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the message processing and security parameters to be used in sending notifications to a particular management target.

The remaining statements are explained separately.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- Defining and Configuring the Trap Target Parameters
- Applying Target Parameters

targets

Syntax

```
targets {
  address;
}
```

Hierarchy Level [edit snmp trap-group *group-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure one or more systems to receive SNMP traps.

Options ***address***—IPv4 or IPv6 address of the system to receive traps. You must specify an address, not a hostname.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- Configuring SNMP Trap Groups

traceoptions

Syntax

```
traceoptions {
  file filename <files number> <match regular-expression> <size size> <world-readable |
    no-world-readable>;
  flag flag;
  no-remote-trace;
}
```

Hierarchy Level [edit snmp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description The output of the tracing operations is placed into log files in the **/var/log** directory. Each log file is named after the SNMP agent that generates it. Currently, the following logs are created in the **/var/log** directory when the **traceoptions** statement is used:

- chassisd
- craftd
- ilmid
- mib2d
- rmopd
- serviced
- snmpd

Options

file *filename*—By default, the name of the log file that records trace output is the name of the process being traced (for example, **mib2d** or **snmpd**). Use this option to specify another name.

files *number*—(Optional) Maximum number of trace files per SNMP subagent. When a trace file (for example, **snmpd**) reaches its maximum size, it is archived by being renamed to **snmpd.0**. The previous **snmpd.1** is renamed to **snmpd.2**, and so on. The oldest archived file is deleted.

Range: 2 through 1000 files

Default: 10 files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements:

- **all**—Log all SNMP events.
- **configuration**—Log reading of configuration at the **[edit snmp]** hierarchy level.
- **database**—Log events involving storage and retrieval in the events database.
- **events**—Log important events.
- **general**—Log general events.

- **interface-stats**—Log physical and logical interface statistics.
- **nonvolatile-sets**—Log nonvolatile SNMP set request handling.
- **pdu**—Log SNMP request and response packets.
- **policy**—Log policy processing.
- **protocol-timeouts**—Log SNMP response timeouts.
- **routing-socket**—Log routing socket calls.
- **server**—Log communication with processes that are generating events.
- **subagent**—Log subagent restarts.
- **timer-events**—Log internally generated events.
- **varbind-error**—Log variable binding errors.

match *regular-expression*—(Optional) Refine the output to include lines that contain the regular expression.

size *size*—(Optional) Maximum size, in kilobytes (KB), of each trace file before it is closed and archived.

Range: 10 KB through 1 GB

Default: 1000 KB

world-readable | no-world-readable—(Optional) By default, log files can be accessed only by the user who configures the tracing operation. The **world-readable** option enables any user to read the file. To explicitly set the default behavior, use the **no-world-readable** option.

Required Privilege Level

snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- Tracing SNMP Activity on a Device Running Junos OS

trap-group

Syntax

```
trap-group group-name {
  categories {
    category;
  }
  destination-port port-number;
  routing-instance instance;
  targets {
    address;
  }
  version (all | v1 | v2);
}
```

Hierarchy Level [edit snmp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Create a named group of hosts to receive the specified trap notifications. The name of the trap group is embedded in SNMP trap notification packets as one variable binding (varbind) known as the community name. At least one trap group must be configured for SNMP traps to be sent.

Options *group-name*—Name of the trap group. If the name includes spaces, enclose it in quotation marks (" ").

The remaining statements are explained separately.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- Configuring SNMP Trap Groups

trap-options

Syntax

```
trap-options {
  agent-address outgoing-interface;
  source-address address;
}
```

Hierarchy Level [edit snmp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Using SNMP trap options, you can set the source address of every SNMP trap packet sent by the router to a single address, regardless of the outgoing interface. In addition, you can set the agent address of each SNMPv1 trap. For more information about the contents of SNMPv1 traps, see RFC 1157.

The remaining statements are explained separately.

Default Disabled

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- Configuring Source and Agent Addresses for SNMP Traps

type

Syntax type (inform | trap);

Hierarchy Level [edit snmp v3 notify *name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the type of notification.

Options **inform**—Defines the type of notification as an inform. SNMP informs are confirmed notifications.

trap—Defines the type of notification as a trap. SNMP traps are unconfirmed notifications.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- Configuring SNMP Informs
- Configuring the SNMPv3 Trap Notification

type

Syntax type *type*;

Hierarchy Level [edit snmp rmon **event** *index*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Type of notification generated when a threshold is crossed.

Options *type*—Type of notification:

- **log**—Add an entry to **logTable**.
- **log-and-trap**—Send an SNMP trap and make a log entry.
- **none**—No notifications are sent.
- **snmptrap**—Send an SNMP trap.

Default: log-and-trap

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- Configuring an Event Entry and Its Attributes

v3

Syntax

```
v3 {
  notify name {
    tag tag-name;
    type trap;
  }
  notify-filter profile-name {
    oid object-identifier (include | exclude);
  }
  snmp-community community-index {
    community-name community-name;
    security-name security-name;
    tag tag-name;
  }
  target-address target-address-name {
    address address;
    address-mask address-mask;
    inform-retry-count number;
    inform-timeout seconds;
    port port-number;
    routing-instance instance;
    tag-list tag-list;
    target-parameters target-parameters-name;
  }
  target-parameters target-parameters-name {
    notify-filter profile-name;
    parameters {
      message-processing-model (v1 | v2c | V3);
      security-level (authentication | none | privacy);
      security-model (usm | v1 | v2c);
      security-name security-name;
    }
  }
  usm {
    local-engine {
      user username {
        authentication-md5 {
          authentication-password authentication-password;
        }
        authentication-sha {
          authentication-password authentication-password;
        }
        authentication-none;
        privacy-aes128 {
          privacy-password privacy-password;
        }
        privacy-des {
          privacy-password privacy-password;
        }
        privacy-des {
          privacy-password privacy-password;
        }
        privacy-none;
```

```
        }
      }
      remote-engine engine-id {
        user username {
          authentication-md5 {
            authentication-password authentication-password;
          }
          authentication-sha {
            authentication-password authentication-password;
          }
          authentication-none;
          privacy-aes128 {
            privacy-password privacy-password;
          }
          privacy-des {
            privacy-password privacy-password;
          }
          privacy-3des {
            privacy-password privacy-password;
          }
          privacy-none {
            privacy-password privacy-password;
          }
        }
      }
    }
    vacm {
      access {
        group group-name {
          default-context-prefix {
            security-model (any | usm | v1 | v2c) {
              security-level (authentication | none | privacy) {
                notify-view view-name;
                read-view view-name;
                write-view view-name;
              }
            }
          }
        }
      }
      security-to-group {
        security-model (usm | v1 | v2c) {
          security-name security-name {
            group group-name;
          }
        }
      }
    }
  }
```

Hierarchy Level [edit snmp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure SNMPv3.

The remaining statements are explained separately.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Minimum SNMPv3 Configuration on a Device Running Junos OS

vacm

Syntax

```
vacm {
  access {
    group group-name {
      default-context-prefix {
        security-model (any | usm | v1 | v2c) {
          security-level (authentication | none | privacy) {
            notify-view view-name;
            read-view view-name;
            write-view view-name;
          }
        }
      }
    }
  }
  security-to-group {
    security-model (usm | v1 | v2c);
    security-name security-name {
      group group-name;
    }
  }
}
```

Hierarchy Level [edit snmp v3]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure view-based access control model (VACM) information.

The remaining statements are explained separately.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Defining Access Privileges for an SNMP Group

variable

Syntax variable *oid-variable*;

Hierarchy Level [edit snmp rmon **alarm** *index*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Object identifier (OID) of MIB variable to be monitored.

Options ***oid-variable***—OID of the MIB variable that is being monitored. The OID can be a dotted decimal (for example, **1.3.6.1.2.1.2.1.2.2.1.10.1**). Alternatively, use the MIB object name (for example, **ifInOctets.1**).

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the Variable

version

Syntax version (all | v1 | v2);

Hierarchy Level [edit snmp **trap-group** *group-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the version number of SNMP traps.

Options **all**—Send an SNMPv1 and SNMPv2 trap for every trap condition.

v1—Send SNMPv1 traps only.

v2—Send SNMPv2 traps only.

Default: all

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring SNMP Trap Groups

view (Configuring a MIB View)

Syntax

```
view view-name {
  oid object-identifier (include | exclude);
}
```

Hierarchy Level [edit snmp]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define a MIB view. A MIB view identifies a group of MIB objects. Each MIB object in a view has a common OID prefix. Each object identifier represents a subtree of the MIB object hierarchy. The **view** statement uses a view to specify a group of MIB objects on which to define access. To enable a view, you must associate the view with a community by including the **view** statement at the **[edit snmp community *community-name*]** hierarchy level.

NOTE: To remove an OID completely, use the delete view all oid oid-number command but omit the include parameter.

Options ***view-name***—Name of the view

The remaining statement is explained separately.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- Configuring MIB Views
- Associating MIB Views with an SNMP User Group
- **community on page 3320**

view (Associating a MIB View with a Community)

Syntax view *view-name*;

Hierarchy Level [edit snmp **community** *community-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Associate a view with a community. A view represents a group of MIB objects.

Options ***view-name***—Name of the view. You must use a view name already configured in the **view** statement at the **[edit snmp]** hierarchy level.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring the SNMP Community String

write-view

Syntax write-view *view-name*;

Hierarchy Level [edit snmp v3 vacm access group *group-name* default-context-prefix security-model (any | usm | v1 | v2c) **security-level** (authentication | none | privacy)]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Associate the view with a community or a group name (SNMPv3).

Options ***view-name***—The name of the view to which the SNMP user group has access.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation
- Configuring MIB Views
- Configuring the Write View

Operational Mode Commands for SNMP

clear snmp rmon history

Syntax **clear snmp rmon history <*interface-name* | all>**

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Delete the samples of Ethernet statistics collected, but do not delete the RMON history configuration.

The **clear snmp rmon history** command deletes all the samples collected for the interface configured for the history group, but not the configuration of that group. If you want to delete the RMON history group configuration, you must use the **delete snmp rmon history** configuration-mode command.

Options *interface-name*—Delete the samples of Ethernet statistics collected for this interface.

all—Delete the samples of Ethernet statistics collected for all interfaces that have been configured for RMON monitoring.

Required Privilege Level clear

Related Documentation

- **show snmp rmon history on page 3392**

clear snmp statistics

Syntax clear snmp statistics

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Clear Simple Network Management Protocol (SNMP) statistics.

Options This command has no options.

Required Privilege Level clear

Related Documentation
- **show snmp statistics on page 3395**

List of Sample Output **clear snmp statistics on page 3372**

Output Fields See **show snmp statistics** for an explanation of output fields.

clear snmp statistics In the following example, SNMP statistics are displayed before and after the **clear snmp statistics** command is issued:

```
user@host> show snmp statistics
SNMP statistics:
  Input:
    Packets: 8, Bad versions: 0, Bad community names: 0,
    Bad community uses: 0, ASN parse errors: 0,
    Too bigs: 0, No such names: 0, Bad values: 0,
    Read onlys: 0, General errors: 0,
    Total request varbinds: 8, Total set varbinds: 0,
    Get requests: 0, Get nexts: 8, Set requests: 0,
    Get responses: 0, Traps: 0,
    Silent drops: 0, Proxy drops 0
  Output:
    Packets: 2298, Too bigs: 0, No such names: 0,
    Bad values: 0, General errors: 0,
    Get requests: 0, Get nexts: 0, Set requests: 0,
    Get responses: 8, Traps: 2290


user@host> clear snmp statistics


user@host> show snmp statistics
SNMP statistics:
  Input:
    Packets: 0, Bad versions: 0, Bad community names: 0,
    Bad community uses: 0, ASN parse errors: 0,
    Too bigs: 0, No such names: 0, Bad values: 0,
    Read onlys: 0, General errors: 0,
    Total request varbinds: 0, Total set varbinds: 0,
    Get requests: 0, Get nexts: 0, Set requests: 0,
    Get responses: 0, Traps: 0,
    Silent drops: 0, Proxy drops 0
  Output:
    Packets: 0, Too bigs: 0, No such names: 0,
```

```
Bad values: 0, General errors: 0,
Get requests: 0, Get nexts: 0, Set requests: 0,
Get responses: 0, Traps: 0
```

request snmp spoof-trap

Syntax

request snmp spoof-trap
<*trap*> variable-bindings <*object*> <*instance*> <*value*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Spoof (mimic) the behavior of a Simple Network Management Protocol (SNMP) trap.

Options

<*trap*>—Name of the trap to spoof.

variable-bindings <*object*> <*instance*> <*value*>—(Optional) List of variables and values to include in the trap. Each variable binding is specified as an object name, the object instance, and the value (for example, **ifIndex[14] = 14**). Enclose the list of variable bindings in quotation marks (“ “) and use a comma to separate each object name, instance, and value definition (for example, **variable-bindings “ifIndex[14] = 14, ifAdminStatus[14] = 1, ifOperStatus[14] = 2”**). Objects included in the trap definition that do not have instances and values specified as part of the command are included in the trap and spoofed with automatically generated instances and values.

<*dummy name*>—A dummy trap name to display the list of available traps.

Question mark (**?**)—Question mark**?** to display possible completions.

Required Privilege Level

request

List of Sample Output

request snmp spoof-trap (with Variable Bindings) on page 3374
request snmp spoof-trap (Illegal Trap Name) on page 3374
request snmp spoof-trap (Question Mark ?) on page 3378

request snmp spoof-trap (with Variable Bindings)

```
user@host> request snmp spoof-trap linkUp variable-bindings “ifIndex[14] = 14, ifAdminStatus[14]
= 1, ifOperStatus[14] = 2”
Spoof trap request result: trap sent successfully
```

request snmp spoof-trap (Illegal Trap Name)

```
user@host> request snmp spoof-trap xx
Spoof trap request result: trap not found

Allowed Traps:
adslAtucInitFailureTrap
  adslAtucPerfESsThreshTrap
  adslAtucPerfLofsThreshTrap
  adslAtucPerfLolsThreshTrap
  adslAtucPerfLossThreshTrap
  adslAtucPerfLprsThreshTrap
  adslAtucRateChangeTrap
  adslAturPerfESsThreshTrap
  adslAturPerfLofsThreshTrap
  adslAturPerfLossThreshTrap
  adslAturPerfLprsThreshTrap
  adslAturRateChangeTrap
  apsEventChannelMismatch
  apsEventFEPLF
  apsEventModeMismatch
  apsEventPSBF
```

```
apsEventSwitchover
authenticationFailure
bfdSessDown
bfdSessUp
bgpBackwardTransition
bgpEstablished
coldStart
dlswTrapCircuitDown
dlswTrapCircuitUp
dlswTrapTConnDown
dlswTrapTConnPartnerReject
dlswTrapTConnProtViolation
dlswTrapTConnUp
dsx1LineStatusChange
dsx3LineStatusChange
entConfigChange
fallingAlarm
frDLCIStatusChange
ggsnTrapChanged
ggsnTrapCleared
ggsnTrapNew
gmplsTunnelDown
ifMauJabberTrap
ipv6IfStateChange
isisAreaMismatch
isisAttemptToExceedMaxSequence
isisAuthenticationFailure
isisAuthenticationTypeFailure
isisCorruptedLSPDetected
isisDatabaseOverload
isisIDLenMismatch
isisLSPTooLargeToPropagate
isisManualAddressDrops
isisMaxAreaAddressesMismatch
isisOriginatingLSPBufferSizeMismatch
isisOwnLSPPurge
isisProtocolsSupportedMismatch
isisRejectedAdjacency
isisSequenceNumberSkip
isisVersionSkew
jnxAccessAuthServerDisabled
jnxAccessAuthServerEnabled
jnxAccessAuthServiceDown
jnxAccessAuthServiceUp
jnxBfdSessDetectionTimeHigh
jnxBfdSessTxIntervalHigh
jnxBgpM2BackwardTransition
jnxBgpM2Established
jnxCmCfgChange
jnxCmRescueChange
jnxCollFlowOverload
jnxCollFlowOverloadCleared
jnxCollFtpSwitchover
jnxCollMemoryAvailable
jnxCollMemoryUnavailable
jnxCollUnavailableDest
jnxCollUnavailableDestCleared
jnxCollUnsuccessfulTransfer
jnxDfcHardMemThresholdExceeded
jnxDfcHardMemUnderThreshold
jnxDfcHardPpsThresholdExceeded
```

```
jnxDfcHardPpsUnderThreshold
jnxDfcSoftMemThresholdExceeded
jnxDfcSoftMemUnderThreshold
jnxDfcSoftPpsThresholdExceeded
jnxDfcSoftPpsUnderThreshold
jnxEventTrap
jnxExampleStartup
jnxFEBSwitchover
jnxFanFailure
jnxFanOK
jnxFruCheck
jnxFruFailed
jnxFruInsertion
jnxFruOK
jnxFruOffline
jnxFruOnline
jnxFruPowerOff
jnxFruPowerOn
jnxFruRemoval
jnxHardDiskFailed
jnxHardDiskMissing
jnxJsAvPatternUpdateTrap
jnxJsChassisClusterSwitchover
jnxJsFwAuthCapacityExceeded
jnxJsFwAuthFailure
jnxJsFwAuthServiceDown
jnxJsFwAuthServiceUp
jnxJsNatAddrPoolThresholdStatus
jnxJsScreenAttack
jnxJsScreenCfgChange
jnxLdpLspDown
jnxLdpLspUp
jnxLdpSesDown
jnxLdpSesUp
jnxMIMstCistPortLoopProtectStateChangeTrap
jnxMIMstCistPortRootProtectStateChangeTrap
jnxMIMstErrTrap
jnxMIMstGenTrap
jnxMIMstInvalidBpduRxdTrap
jnxMIMstMstiPortLoopProtectStateChangeTrap
jnxMIMstMstiPortRootProtectStateChangeTrap
jnxMIMstNewRootTrap
jnxMIMstProtocolMigrationTrap
jnxMIMstRegionConfigChangeTrap
jnxMIMstTopologyChgTrap
jnxMacChangedNotification
jnxMplsLdpInitSesThresholdExceeded
jnxMplsLdpPathVectorLimitMismatch
jnxMplsLdpSessionDown
jnxMplsLdpSessionUp
jnxOspfv3IfConfigError
jnxOspfv3IfRxBadPacket
jnxOspfv3IfStateChange
jnxOspfv3LsdbApproachingOverflow
jnxOspfv3LsdbOverflow
jnxOspfv3NbrRestartHelperStatusChange
jnxOspfv3NbrStateChange
jnxOspfv3NssaTranslatorStatusChange
jnxOspfv3RestartStatusChange
jnxOspfv3VirtIfConfigError
jnxOspfv3VirtIfRxBadPacket
```

```
jnxOspfv3VirtIfStateChange
jnxOspfv3VirtNbrRestartHelperStatusChange
jnxOspfv3VirtNbrStateChange
jnxOtnAlarmCleared
jnxOtnAlarmSet
jnxOverTemperature
jnxPMonOverloadCleared
jnxPMonOverloadSet
jnxPingEgressJitterThresholdExceeded
jnxPingEgressStdDevThresholdExceeded
jnxPingEgressThresholdExceeded
jnxPingIngressJitterThresholdExceeded
jnxPingIngressStddevThresholdExceeded
jnxPingIngressThresholdExceeded
jnxPingRttJitterThresholdExceeded
jnxPingRttStdDevThresholdExceeded
jnxPingRttThresholdExceeded
jnxPortBpduErrorStatusChangeTrap
jnxPortLoopProtectStateChangeTrap
jnxPortRootProtectStateChangeTrap
jnxPowerSupplyFailure
jnxPowerSupplyOK
jnxRedundancySwitchover
jnxRmonAlarmGetFailure
jnxRmonGetOk
jnxSecAccessIfMacLimitExceeded
jnxSecAccessdsRateLimitCrossed
jnxSonetAlarmCleared
jnxSonetAlarmSet
jnxSpSvcSetCpuExceeded
jnxSpSvcSetCpuOk
jnxSpSvcSetZoneEntered
jnxSpSvcSetZoneExited
jnxStormEventNotification
jnxSyslogTrap
jnxTemperatureOK
jnxVccpPortDown
jnxVccpPortUp
jnxVpnIfDown
jnxVpnIfUp
jnxVpnPwDown
jnxVpnPwUp
jnxl2aldGlobalMacLimit
jnxl2aldInterfaceMacLimit
jnxl2aldRoutingInstMacLimit
linkDown
linkUp
lldpRemTablesChange
mfrMibTrapBundleLinkMismatch
mplsLspChange
mplsLspDown
mplsLspInfoChange
mplsLspInfoDown
mplsLspInfoPathDown
mplsLspInfoPathUp
mplsLspInfoUp
mplsLspPathDown
mplsLspPathUp
mplsLspUp
mplsNumVrfRouteMaxThreshExceeded
mplsNumVrfRouteMidThreshExceeded
```

```
  mplsNumVrfSecIllglLblThrshExcd
  mplsTunnelDown
  mplsTunnelReoptimized
  mplsTunnelRerouted
  mplsTunnelUp
  mplsVrfIfDown
  mplsVrfIfUp
  mplsXCDown
  mplsXCUp
  msdpBackwardTransition
  msdpEstablished
  newRoot
  ospfIfAuthFailure
  ospfIfConfigError
  ospfIfRxBadPacket
  ospfIfStateChange
  ospfLsdbApproachingOverflow
  ospfLsdbOverflow
  ospfMaxAgeLsa
  ospfNbrStateChange
  ospfOriginateLsa
  ospfTxRetransmit
  ospfVirtIfAuthFailure
  ospfVirtIfConfigError
  ospfVirtIfRxBadPacket
  ospfVirtIfStateChange
  ospfVirtIfTxRetransmit
  ospfVirtNbrStateChange
  pethMainPowerUsageOffNotification
  pethMainPowerUsageOnNotification
  pethPsePortOnOffNotification
  pingProbeFailed
  pingTestCompleted
  pingTestFailed
  ptopoConfigChange
  risingAlarm
  rpMauJabberTrap
  sdlcLSStatusChange
  sdlcPortStatusChange
  topologyChange
  traceRoutePathChange
  traceRouteTestCompleted
  traceRouteTestFailed
  vrrpTrapAuthFailure
  vrrpTrapNewMaster
  warmStart
```

request snmp spoof-trap (Question Mark ?)

```
user@host> request snmp spoof-trap ?
Possible completions:
  <trap>               The name of the trap to spoof
  adslAtucInitFailureTrap
  adslAtucPerfESsThreshTrap
  adslAtucPerfLofsThreshTrap
  adslAtucPerfLolsThreshTrap
  adslAtucPerfLossThreshTrap
  adslAtucPerfLprsThreshTrap
  adslAtucRateChangeTrap
  adslAturPerfESsThreshTrap
  adslAturPerfLofsThreshTrap
  adslAturPerfLossThreshTrap
  adslAturPerfLprsThreshTrap
```

```
  adslAturRateChangeTrap
  apsEventChannelMismatch
  apsEventFEPLF
  apsEventModeMismatch
  apsEventPSBF
  apsEventSwitchover
  authenticationFailure
  bfdSessDown
  bfdSessUp
  bgpBackwardTransition
  bgpEstablished
  coldStart
  dlswTrapCircuitDown
  dlswTrapCircuitUp
---(more 10%)---
```

show snmp health-monitor

Syntax show snmp health-monitor
<alarms <detail>> | <logs>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display information about Simple Network Management Protocol (SNMP) health monitor alarms and logs.

Options none—Display information about all health monitor alarms and logs.

alarms <detail>—(Optional) Display detailed information about health monitor alarms.

logs—(Optional) Display information about health monitor logs.

Required Privilege Level view

List of Sample Output **show snmp health-monitor on page 3382**
show snmp health-monitor alarms detail on page 3384

Output Fields Table 444 on page 3380 describes the output fields for the **show snmp health-monitor** command. Output fields are listed in the approximate order in which they appear.

Table 444: show snmp health-monitor Output Fields

Field Name	Field Description	Level of Output
Alarm Index	Alarm identifier.	All levels
Variable description	Description of the health monitor object instance being monitored.	All levels
Variable name	Name of the health monitor object instance being monitored.	All levels
Value	Current value of the monitored variable in the most recent sample interval.	All levels

Table 444: show snmp health-monitor Output Fields *(continued)*

Field Name	Field Description	Level of Output

State	State of the alarm or event entry:
	• Alarms:
	  • **active**—Entry is fully configured and activated.
	  • **falling threshold crossed**—Value of the variable has crossed the lower threshold limit.
	  • **rising threshold crossed**—Value of the variable has crossed the upper threshold limit.
	  • **under creation**—Entry is being configured and is not yet activated.
	  • **startup**—Alarm is waiting for the first sample of the monitored variable.
	  • **object not available**—Monitored variable of that type is not available to the health monitor agent.
	  • **instance not available**—Monitored variable's instance is not available to the health monitor agent.
	  • **object type invalid**—Monitored variable is not a numeric value.
	  • **object processing errored**—An error occurred when the monitored variable was processed.
	  • **unknown**—State is not one of the above.
		All levels

Variable OID	Object ID to which the variable name is resolved. The format is **x.x.x.x**.	**detail**
Sample type	Method of sampling the monitored variable and calculating the value to compare against the upper and lower thresholds. It can have the value of **absolute value** or **delta value**.	**detail**

Startup alarm	Alarm that might be sent when this entry is first activated, depending on the following criteria:
	• Alarm is sent when one of the following situations exists:
	  • Value of the alarm is above or equal to the rising threshold and the startup type is either **rising alarm** or **rising or falling alarm.**
	  • Value of the alarm is below or equal to the falling threshold and the startup type is either **falling alarm** or **rising or falling alarm**.
	• Alarm is *not* sent when one of the following situations exists:
	  • Value of the alarm is above or equal to the rising threshold and the startup type is **falling alarm**.
	  • Value of the alarm is below or equal to the falling threshold and the startup type is **rising alarm**.
	  • Value of the alarm is between the thresholds.
		detail

Owner	Name of the entry configured by the user. If the entry was created through the CLI, the owner has **monitor** prepended to it.	**detail**
Creator	Mechanism by which the entry was configured (**Health Monitor**).	**detail**
Sample interval	Time period between samples (in seconds).	**detail**
Rising threshold	Upper limit threshold value as a percentage of the maximum possible value.	**detail**

Table 444: show snmp health-monitor Output Fields *(continued)*

Field Name	Field Description	Level of Output
Falling threshold	Lower limit threshold value as a percentage of the maximum possible value.	**detail**
Rising event index	Event triggered when the rising threshold is crossed.	**detail**
Falling event index	Event triggered when the falling threshold is crossed.	**detail**

show snmp health-monitor

```
user@host> show snmp health-monitor

Alarm
Index  Variable description                          Value State

 32768 Health Monitor: root file system utilization
       jnxHrStoragePercentUsed.1                       58 active

 32769 Health Monitor: /config file system utilization
       jnxHrStoragePercentUsed.2                        0 active

 32770 Health Monitor: RE 0 CPU utilization
       jnxOperatingCPU.9.1.0.0                          0 active

 32773 Health Monitor: RE 0 Memory utilization
       jnxOperatingBuffer.9.1.0.0                      35 active

 32775 Health Monitor: jkernel daemon CPU utilization
       Init daemon                                      0 active
       Chassis daemon                                  50 active
       Firewall daemon                                  0 active
       Interface daemon                                 5 active
       SNMP daemon                                     11 active
       MIB2 daemon                                     42 active
       Sonet APS daemon                                 0 active
       VRRP daemon                                      0 active
       Alarm daemon                                     3 active
       PFE daemon                                       0 active
       CRAFT daemon                                     0 active
       Traffic sampling control daemon                  0 active
       Ilmi daemon                                      0 active
       Remote operations daemon                         0 active
       CoS daemon                                       0 active
       Pic Services Logging daemon                      0 active
       Internal Routing Service Daemon                  3 active
       Network Access Service daemon                    0 active
       Forwarding UDP daemon                            0 active
       Routing socket proxy daemon                      0 active
       Disk Monitoring daemon                           1 active
       Inet daemon                                      0 active
       Syslog daemon                                    0 active
       Adaptive Services PIC daemon                     0 active
       ECC parity errors logging Daemon                 0 active
       Layer 2 Tunneling Protocol daemon                0 active
       PPPoE daemon                                     3 active
       Redundancy device daemon                         0 active
       PPP daemon                                       0 active
       Dynamic Flow Capture Daemon                      0 active
```

```
32776 Health Monitor: jroute daemon CPU utilization
      Routing protocol daemon                             1 active
      Management daemon                                   0 active
      Management daemon                                   0 active
      Command line interface                              4 active
      Periodic Packet Management daemon                   0 active
      Link Management daemon                              0 active
      Pragmatic General Multicast daemon                  0 active
      Bidirectional Forwarding Detection daemon           0 active
      SRC daemon                                          0 active
      audit daemon                                        0 active
      Event daemon                                        0 active

32777 Health Monitor: jcrypto daemon CPU utilization
      IPSec Key Management daemon                         0 active

32779 Health Monitor: jkernel daemon Memory utilization
      Init daemon                                     47384 active
      Chassis daemon                                  20204 active
      Firewall daemon                                  1956 active
      Interface daemon                                 3340 active
      SNMP daemon                                      4540 active
      MIB2 daemon                                      3880 active
      Sonet APS daemon                                 2632 active
      VRRP daemon                                      2672 active
      Alarm daemon                                     1856 active
      PFE daemon                                       2600 active
      CRAFT daemon                                     2000 active
      Traffic sampling control daemon                  3164 active
      Ilmi daemon                                      2132 active
      Remote operations daemon                         2964 active
      CoS daemon                                       3044 active
      Pic Services Logging daemon                      1944 active
      Internal Routing Service Daemon                  1392 active
      Network Access Service daemon                    1992 active
      Forwarding UDP daemon                            1876 active
      Routing socket proxy daemon                      1296 active
      Disk Monitoring daemon                           1180 active
      Inet daemon                                      1296 active
      Syslog daemon                                    1180 active
      Adaptive Services PIC daemon                     3220 active
      ECC parity errors logging Daemon                 1100 active
      Layer 2 Tunneling Protocol daemon                3372 active
      PPPoE daemon                                     1424 active
      Redundancy device daemon                         1820 active
      PPP daemon                                       2060 active
      Dynamic Flow Capture Daemon                     10740 active
32780 Health Monitor: jroute daemon Memory utilization
      Routing protocol daemon                          8104 active
      Management daemon                               13360 active
      Management daemon                               19252 active
      Command line interface                           9912 active
      Periodic Packet Management daemon                1484 active
      Link Management daemon                           2016 active
      Pragmatic General Multicast daemon               1968 active
      Bidirectional Forwarding Detection daemon        1956 active
      SRC daemon                                       1772 active
      audit daemon                                     1772 active
      Event daemon                                     1808 active
```

```
 32781 Health Monitor: jcrypto daemon Memory utilization
       IPSec Key Management daemon                    5600 active
```

show snmp health-monitor alarms detail

```
user@host> show snmp health-monitor alarms detail

Alarm Index 32768:
  Variable name                       jnxHrStoragePercentUsed.1
  Variable OID                        1.3.6.1.4.1.2636.3.31.1.1.1.1.1
  Sample type                         absolute value
  Startup alarm                       rising alarm
  Owner                               Health Monitor: root file system
                                      utilization
  Creator                             Health Monitor
  State                               active
  Sample interval                 300 seconds
  Rising threshold                 80
  Falling threshold                70
  Rising event index            32768
  Falling event index           32768
    Instance Value: 58
    Instance State: active

Alarm Index 32769:
  Variable name                       jnxHrStoragePercentUsed.2
  Variable OID                        1.3.6.1.4.1.2636.3.31.1.1.1.1.2
  Sample type                         absolute value
  Startup alarm                       rising alarm
  Owner                               Health Monitor: /config file system
                                      utilization
  Creator                             Health Monitor
  State                               active
  Sample interval                 300 seconds
  Rising threshold                 80
  Falling threshold                70
  Rising event index            32768
  Falling event index           32768
    Instance Value: 0
    Instance State: active

Alarm Index 32770:
  Variable name                       jnxOperatingCPU.9.1.0.0
  Variable OID                        1.3.6.1.4.1.2636.3.1.13.1.8.9.1.0.0
  Sample type                         absolute value
  Startup alarm                       rising alarm
  Owner                              Health Monitor: RE 0 CPU utilization


  Creator                             Health Monitor
  State                               active
  Sample interval                 300 seconds
  Rising threshold                 80
  Falling threshold                70
  Rising event index            32768
  Falling event index           32768
    Instance Value: 0
    Instance State: active

Alarm Index 32773:
  Variable name                       jnxOperatingBuffer.9.1.0.0
  Variable OID                        1.3.6.1.4.1.2636.3.1.13.1.11.9.1.0.0
  Sample type                         absolute value
```

```
  Startup alarm                       rising alarm
  Owner                               Health Monitor: RE 0 Memory utilization


  Creator                             Health Monitor
  State                               active
  Sample interval                 300 seconds
  Rising threshold                 80
  Falling threshold                70
  Rising event index            32768
  Falling event index           32768
    Instance Value: 35
    Instance State: active

Alarm Index 32775:
  Variable name                       sysApplElmtRunCPU.3
  Variable OID                        1.3.6.1.2.1.54.1.2.3.1.9.3
  Sample type                         delta value
  Startup alarm                       rising alarm
  Owner                               Health Monitor: jkernel daemon CPU
                                      utilization
  Creator                             Health Monitor
  State                               active
  Sample interval                 300 seconds
  Rising threshold              24000
  Falling threshold             21000
  Rising event index            32768
  Falling event index           32768
    Instance Name: sysApplElmtRunCPU.3.1.1
    Instance Description: Init daemon
    Instance Value: 0
    Instance State: active

    Instance Name: sysApplElmtRunCPU.3.2.2786
    Instance Description: Chassis daemon
    Instance Value: 50
    Instance State: active

    Instance Name: sysApplElmtRunCPU.3.3.2938
    Instance Description: Firewall daemon
    Instance Value: 0
    Instance State: active

    Instance Name: sysApplElmtRunCPU.3.4.2942
    Instance Description: Interface daemon
    Instance Value: 5
    Instance State: active

    Instance Name: sysApplElmtRunCPU.3.7.7332
    Instance Description: SNMP daemon
    Instance Value: 11
    Instance State: active

    Instance Name: sysApplElmtRunCPU.3.9.2914
    Instance Description: MIB2 daemon
    Instance Value: 42
    Instance State: active

    Instance Name: sysApplElmtRunCPU.3.12.2916
    Instance Description: Sonet APS daemon
    Instance Value: 0
```

```
    Instance State: active

    Instance Name: sysApplElmtRunCPU.3.13.2917
    Instance Description: VRRP daemon
    Instance Value: 0
    Instance State: active

    Instance Name: sysApplElmtRunCPU.3.14.2787
    Instance Description: Alarm daemon
    Instance Value: 3
    Instance State: active

    Instance Name: sysApplElmtRunCPU.3.15.2940
    Instance Description: PFE daemon
    Instance Value: 0
    Instance State: active

    Instance Name: sysApplElmtRunCPU.3.16.2788
    Instance Description: CRAFT daemon
    Instance Value: 0
    Instance State: active

    Instance Name: sysApplElmtRunCPU.3.17.2918
    Instance Description: Traffic sampling control daemon
---(more 23%)---
```

show snmp inform-statistics

Syntax show snmp inform-statistics

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display information about Simple Network Management Protocol (SNMP) inform requests.

Options This command has no options.

Required Privilege Level view

List of Sample Output **show snmp inform-statistics on page 3387**

Output Fields Table 445 on page 3387 describes the output fields for the **show snmp inform-statistics** command. Output fields are listed in the approximate order in which they appear.

Table 445: show snmp inform-statistics Output Fields

Field Name	Field Description
Target Name	Name of the device configured to receive and respond to SNMP informs.
Address	IP address of the target device.
Sent	Number of informs sent to the target device and acknowledged by the target device.
Pending	Number of informs held in memory pending a response from the target device.
Discarded	Number of informs discarded after the specified number of retransmissions to the target device were attempted.
Timeouts	Number of informs that did not receive an acknowledgement from the target device within the timeout specified.
Probe Failures	Connection failures that occurred (for example, when the target server returned invalid content or you incorrectly configured the target address).

show snmp inform-statistics

```
user@host> show snmp inform-statistics
Inform Request Statistics:
 Target Name: TA1_v3_md5_none Address: 172.17.20.184
  Sent: 176, Pending: 0
  Discarded: 0, Timeouts: 0, Probe Failures: 0
 Target Name: TA2_v3_sha_none Address: 192.168.110.59
  Sent: 0, Pending: 4
  Discarded: 84, Timeouts: 0, Probe Failures: 258
 Target Name: TA5_v2_none Address: 172.17.20.184
  Sent: 0, Pending: 0
  Discarded: 2, Timeouts: 10, Probe Failures: 0
```

show snmp rmon

Syntax show snmp rmon
<alarms <brief | detail> | events <brief | detail> | logs>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display information about Simple Network Management Protocol (SNMP) Remote Monitoring (RMON) alarms and events.

Options none—Display information about all RMON alarms and events.

alarms—(Optional) Display information about RMON alarms.

brief | detail—(Optional) Display brief or detailed information about RMON alarms or events.

events—(Optional) Display information about RMON events.

logs—(Optional) Display information about RMON monitoring logs.

Required Privilege Level view

List of Sample Output **show snmp rmon on page 3390**
show snmp rmon alarms detail on page 3390
show snmp rmon events detail on page 3391

Output Fields Table 446 on page 3388 describes the output fields for the **show snmp rmon** command. Output fields are listed in the approximate order in which they appear.

Table 446: show snmp rmon Output Fields

Field Name	Field Description	Level of Output
Alarm Index	Alarm identifier.	All levels

Table 446: show snmp rmon Output Fields *(continued)*

Field Name	Field Description	Level of Output
State	State of the alarm or event entry: Alarms: • **active**—Entry is fully configured and activated. • **falling threshold crossed**—Value of the variable has crossed the lower threshold limit. • **rising threshold crossed**—Value of the variable has crossed the upper threshold limit. • **under creation**—Entry is being configured and is not yet activated. • **startup**—Alarm is waiting for the first sample of the monitored variable. • **object not available**—Monitored variable of that type is not available to the SNMP agent. • **instance not available**—Monitored variable's instance is not available to the SNMP agent. • **object type invalid**—Monitored variable is not a numeric value. • **object processing errored**—An error occurred when the monitored variable was processed. • **unknown**—State is not one of the above. Events: • **active**—Entry has been fully configured and activated. • **under creation**—Entry is being configured and is not yet activated. • **unknown**—State is not one of the above.	All levels
Variable name	Name of the SNMP object instance being monitored.	All levels
Event Index	Event identifier.	All levels
Type	Type of notification made when an event is triggered. It can be one of the following: • **log**—A system log message is generated and an entry is made to the log table. • **snmptrap**—An SNMP trap is sent to the configured destination. • **log and trap**—A system log message is generated, an entry is made to the log table, and an SNMP trap is sent to the configured destination. • **none**—Neither log nor trap will be sent.	**detail**
Last Event	Date and time of the last event. It has the format *yyyy-mm-dd hh:mm:ss timezone*.	**brief**
Community	Identifies the trap group used for sending the SNMP trap.	**detail**
Variable OID	Object ID to which the variable name is resolved. The format is **x.x.x.x**.	**detail**
Sample type	Method of sampling the monitored variable and calculating the value to compare against the upper and lower thresholds. It can have the value of **absolute value** or **delta value**.	**detail**

Table 446: show snmp rmon Output Fields *(continued)*

Field Name	Field Description	Level of Output

Startup alarm	Alarm that might be sent when this entry is first activated, depending on the following criteria:
	• Alarm is sent when one of the following situations exists:
	  ◦ Value of the alarm is above or equal to the rising threshold and the startup type is either **rising alarm** or **rising or falling alarm.**
	  ◦ Value of the alarm is below or equal to the falling threshold and the startup type is either **falling alarm** or **rising or falling alarm**.
	• Alarm is *not* sent when one of the following situations exists:
	  ◦ Value of the alarm is above or equal to the rising threshold and the startup type is **falling alarm**.
	  ◦ Value of the alarm is below or equal to the falling threshold and the startup type is **rising alarm**.
	  ◦ Value of the alarm is between the thresholds.
		detail

Owner	Name of the entry configured by the user. If the entry was created through the CLI, the owner has **monitor** prepended to it.	**detail**
Creator	Mechanism by which the entry was configured (**CLI** or **SNMP**).	**detail**
Sample interval	Time period between samples (in seconds).	**detail**
Rising threshold	Upper limit threshold value configured by the user.	**detail**
Falling threshold	Lower limit threshold value configured by the user.	**detail**
Rising event index	Event triggered when the rising threshold is crossed.	**detail**
Falling event index	Event triggered when the falling threshold is crossed.	**detail**
Current value	Current value of the monitored variable in the most recent sample interval.	**detail**

show snmp rmon

```
user@host> show snmp rmon
Alarm
Index  State                       Variable name
    1  falling threshold crossed   ifInOctets.1

Event
Index  Type                        Last Event
    1  log and trap                2002-01-30 01:13:01 PST
```

show snmp rmon alarms detail

```
user@host> show snmp rmon alarms detail

Alarm Index 1:
  Variable name                     ifInOctets.1
  Variable OID                      1.3.6.1.2.1.2.2.1.10.1
  Sample type                       delta value
  Startup alarm                     rising or falling alarm
```

```
  Owner                                monitor
  Creator                              CLI
  State                                falling threshold crossed
  Sample interval                   60 seconds
  Rising threshold              100000
  Falling threshold              80000
  Rising event index                 1
  Falling event index                1
  Current value                      0
```

show snmp rmon events detail

```
user@host> show snmp rmon events detail
Event Index 1:
  Type                                 log and trap
  Community                            boy-elroy
  Last event                           2002-01-30 01:13:01 PST
  Creator                              CLI
  State                                active
```

show snmp rmon history

Syntax
show snmp rmon history
<history-index>
<sample-index>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display the contents of the RMON history group.

Options
none—Display all the entries in the RMON history group.

history-index—(Optional) Display the contents of the specified entry in the RMON history group.

sample-index—(Optional) Display the statistics collected for the specified sample within the specified entry in the RMON history group.

Required Privilege Level
view

Related Documentation
- **clear snmp rmon history on page 3371**

List of Sample Output
show snmp rmon history 1 on page 3393
show snmp rmon history 1 sample 15 on page 3394

Output Fields
Table 447 on page 3392 lists the output fields for the **show smp rmon history** command. Output fields are listed in the approximate order in which they appear.

Table 447: show smp rmon history Output Fields

Field Name	Field Description
History Index	Identifies this RMON history entry within the RMON history group.
Owner	The entity that configured this entry. Range is 0 to 32 alphanumeric characters.
Status	The status of the RMON history entry.
Interface or **Data Source**	The ifndex object that identifies the interface that is being monitored.
Interval	The interval (in seconds) configured for this RMON history entry.
Buckets Requested	The requested number of buckets (**intervals**) configured for this RMON history entry.
Buckets Granted	The number of buckets granted for this RMON history entry.

Table 447: show smp rmon history Output Fields *(continued)*

Field Name	Field Description
Sample Index	The sample statistics taken at the specified interval. • **Drop Events**—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • **Octets**—Total number of octets and packets. • **Packets**—Total number of packets. • **Broadcast Packets**—Number of broadcast packets. • **Multicast Packets**—Number of multicast packets. • **CRC errors**—Total number of packets received that had a length (excluding framing bits, but including FCS octets) of between 64 and 1518 octets, inclusive, and had either a bad FCS with an integral number of octets (FCS error) or a bad FCS with a nonintegral number of octets (alignment error). • **Undersize Pkts**—Number of packets received during this sampling interval that were less than 64 octets long (excluding framing bits but including FCS octets) and were otherwise well formed. • **Oversize Pkts**—Number of packets received during the sampling interval that were longer than 1518 octets (excluding framing bits, but including FCS octets) but were otherwise well formed. • **Fragments**—Total number of packets that were less than 64 octets in length (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. Fragment frames normally increment because both runts (which are normal occurrences caused by collisions) and noise hits are counted. • **Jabbers**—Number of frames that were longer than 1518 octets (excluding framing bits, but including FCS octets), and had either an FCS error or an alignment error. This definition of jabber is different from the definition in IEEE-802.3 section 8.2.1.5 (10BASE5) and section 10.3.1.4 (10BASE2). These documents define jabber as the condition in which any packet exceeds 20 ms. The allowed range to detect jabber is from 20 ms to 150 ms. • **Collisions**—Number of Ethernet collisions. • **Utilization(%)**—The best estimate of the mean physical layer network utilization on this interface during this sampling interval, in hundredths of a percent.

show snmp rmon history 1

```
user@host> show snmp rmon history 1
 History Index 1:
 Interface                              171
 Requested Buckets                      50
 Interval                               10

Sample Index 1: Interval Start: Tue Feb 12 04:12:32 2008
    Drop Events         0
    Octets            486
    Packets             2
    Broadcast Packet    0
    Multicast Packets   2
    CRC errors          0
    Undersize Pkts      0
    Oversize Pkts       0
    Fragments           0
    Jabbers             0
```

```
    Collisions          0
    Utilization(%)      0

Sample Index 2: Interval Start: Tue Feb 12 04:12:42 2008
    Drop Events         0
    Octets            486
    Packets             2
    Broadcast Packet    0
    Multicast Packets   2
    CRC errors          0
    Undersize Pkts      0
    Oversize Pkts       0
    Fragments           0
    Jabbers             0
    Collisions          0
    Utilization(%)      0

Sample Index 3: Interval Start: Tue Feb 12 04:12:52 2008
    Drop Events         0
    Octets            486
    Packets             2
    Broadcast Packet    0
    Multicast Packets   2
    CRC errors          0
    Undersize Pkts      0
    Oversize Pkts       0
    Fragments           0
    Jabbers             0
    Collisions          0
    Utilization(%)      0
```

show snmp rmon history 1 sample 15

```
user@host> show snmp rmon history 1 sample 15
 Index 1
Owner    = monitor
Status    = valid
Data Source  = ifIndex.17
Interval   = 1800
Buckets Requested = 50
Buckets Granted = 50


Sample Index 44: Interval Start: Thu Jan  1 00:08:35 1970
Drop Events   = 0
Octetes   = 0
Packets   = 0
Broadcast Pkts  = 0
Multicast Pkts  = 0
CRC Errors  = 0
Undersize Pkts  = 0
Oversize Pkts  = 0
Fragments  = 0
Jabbers   = 0
Collisions  = 0
Utilization (%)  = 0
```

show snmp statistics

Syntax show snmp statistics

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display statistics about Simple Network Management Protocol (SNMP) packets sent and received by the router or switch.

Options This command has no options.

Required Privilege Level view

Related Documentation
- **clear snmp statistics on page 3372**

List of Sample Output **show snmp statistics on page 3398**

Output Fields Table 448 on page 3395 describes the output fields for the **show snmp statistics** command. Output fields are listed in the approximate order in which they appear.

Table 448: show snmp statistics Output Fields

Field Name	Field Description
Input	Information about received packets: • **Packets(snmpInPkts)**—Total number of messages delivered to the SNMP entity from the transport service. • **Bad versions—(snmpInBadVersions**) Total number of messages delivered to the SNMP entity that were for an unsupported SNMP version. • **Bad community names—(snmpInBadCommunityNames)** Total number of messages delivered to the SNMP entity that used an SNMP community name not known to the entity. • **Bad community uses—(snmpInBadCommunityUses)** Total number of messages delivered to the SNMP entity that represented an SNMP operation that was not allowed by the SNMP community named in the message. • **ASN parse errors—(snmpInASNParseErrs)** Total number of ASN.1 or BER errors encountered by the SNMP entity when decoding received SNMP messages. • **Too bigs—(snmpInTooBigs)** Total number of SNMP PDUs delivered to the SNMP entity with an error status field of **tooBig**. • **No such names—(snmpInNoSuchNames)**.Total number of SNMP PDUs delivered to the SNMP entity with an error status field of **noSuchName**. • **Bad values—(snmpInBadValues)** Total number of SNMP PDUs delivered to the SNMP entity with an error status field of **badValue**. • **Read onlys—(snmpInReadOnlys)** Total number of valid SNMP PDUs delivered to the SNMP entity with an error status field of **readOnly**. Only incorrect implementations of SNMP generate this error.

Table 448: show snmp statistics Output Fields *(continued)*

Field Name	Field Description
Input (continued)	• **General errors—(snmpInGenErrs)** Total number of SNMP PDUs delivered to the SNMP entity with an error status field of**genErr**. • **Total requests varbinds—(snmpInTotalReqVars)** Total number of MIB objects retrieved successfully by the SNMP entity as a result of receiving valid SNMP **GetRequest** and **GetNext** PDUs. • **Total set varbinds—(snmpInSetVars)** Total number of MIB objects modified successfully by the SNMP entity as a result of receiving valid SNMP **SetRequest** PDUs. • **Get requests—(snmpInGetRequests)** Total number of SNMP**GetRequest** PDUs that have been accepted and processed by the SNMP entity. • **Get nexts—(snmpInGetNexts)** Total number of SNMP **GetNext** PDUs that have been accepted and processed by the SNMP entity. • **Set requests—(snmpInSetRequests)** Total number of SNMP **SetRequest** PDUs that have been accepted and processed by the SNMP entity. • **Get responses—(snmpInGetResponses)** Total number of SNMP **GetResponse** PDUs that have been accepted and processed by the SNMP entity. • **Traps—(snmpInTraps)** Total number of SNMP traps generated by the SNMP entity. • **Silent drops—(snmpSilentDrops)** Total number of **GetRequest**, **GetNextRequest**, **GetBulkRequest**, **SetRequests**, and **InformRequest** PDUs delivered to the SNMP entity that were silently dropped because the size of a reply containing an alternate response PDU with an empty variable-bindings field was greater than either a local constraint or the maximum message size associated with the originator of the requests. • **Proxy drops.—(snmpProxyDrops)** Total number of **GetRequest**, **GetNextRequest**, **GetBulkRequest**, **SetRequests**, and **InformRequest** PDUs delivered to the SNMP entity that were silently dropped because the transmission of the message to a proxy target failed in such a way (other than a timeout) that no response PDU could be returned. • **Commit pending drops**—Number of SNMP packets for**Set** requests dropped because of a previous pending SNMP **Set** request on the committed configuration. • **Throttle drops**—Number of SNMP packets for any requests dropped reaching the throttle limit.

Table 448: show snmp statistics Output Fields *(continued)*

Field Name	Field Description
V3 Input	Information about SNMP version 3 packets: • **Unknown security models—(snmpUnknownSecurityModels)** Total number of packets received by the SNMP engine that were dropped because they referenced a security model that was not known to or supported by the SNMP engine. • **Invalid messages—(snmpInvalidMsgs)** Number of packets received by the SNMP engine that were dropped because there were invalid or inconsistent components in the SNMP message. • **Unknown pdu handlers—(snmpUnknownPDUHandlers)** Number of packets received by the SNMP engine that were dropped because the PDU contained in the packet could not be passed to an application responsible for handling the PDU type. • **Unavailable contexts—(snmpUnavailableContexts)** Number of requests received for a context that is known to the SNMP engine, but is currently unavailable. • **Unknown contexts—(snmpUnknownContexts)** Total number of requests received for a context that is unknown to the SNMP engine. • **Unsupported security levels—(usmStatsUnsupportedSecLevels)** Total number of packets received by the SNMP engine which were dropped because they requested a security level unknown to the SNMP engine (or otherwise unavailable). • **Not in time windows—(usmStatsNotInTimeWindows)** Total number of packets received by the SNMP engine that were dropped because they appeared outside of the authoritative SNMP engine's window. • **Unknown user names—(usmStatsUnknownUserNames)** Total number of packets received by the SNMP engine that were dropped because they referenced a user that was not known to the SNMP engine. • **Unknown engine ids—(usmStatsUnknownEngineIDs)** Total number of packets received by the SNMP engine that were dropped because they referenced an SNMP engine ID that was not known to the SNMP engine. • **Wrong digests—(usmStatsWrongDigests)** Total number of packets received by the SNMP engine that were dropped because they didn't contain the expected digest value. • **Decryption errors—(usmStatsDecryptionErrors)** Total number of packets received by the SNMP engine that were dropped because they could not be decrypted.

Table 448: show snmp statistics Output Fields *(continued)*

Field Name	Field Description
Output	Information about transmitted packets: • **Packets—(snmpOutPkts)** Total number of messages passed from the SNMP entity to the transport service. • **Too bigs—(snmpOutTooBigs)** Total number of SNMP PDUs generated by the SNMP entity with an error status field of **tooBig**. • **No such names—(snmpOutNoSuchNames)** Total number of SNMP PDUs delivered to the SNMP entity with an error status field of **noSuchName**. • **Bad values—(snmpOutBadValues)** Total number of SNMP PDUs generated by the SNMP entity with an error status field of **badValue**. • **General errors—(snmpOutGenErrs)** Total number of SNMP PDUs generated the SNMP entity with an error status field of **genErr**. • **Get requests—(snmpOutGetRequests)** Total number of SNMP **GetRequest** PDUs generated by the SNMP entity. • **Get nexts—(snmpOutGetNexts)** Total number of SNMP **GetNext** PDUs generated by the SNMP entity. • **Set requests—(snmpOutSetRequests)** Total number of SNMP **SetRequest** PDUs generated by the SNMP entity. • **Get responses—(snmpOutGetResponses)** Total number of SNMP **GetResponse** PDUs generated by the SNMP entity. • **Traps—(snmpOutTraps)** Total number of SNMP traps generated by the SNMP entity.

show snmp statistics

```
user@host> show snmp statistics
SNMP statistics:
  Input:
    Packets: 246213, Bad versions: 12, Bad community names: 12,
    Bad community uses: 0, ASN parse errors: 96,
    Too bigs: 0, No such names: 0, Bad values: 0,
    Read onlys: 0, General errors: 0,
    Total request varbinds: 227084, Total set varbinds: 67,
    Get requests: 44942, Get nexts: 190371, Set requests: 10712,
    Get responses: 0, Traps: 0,
    Silent drops: 0, Proxy drops: 0, Commit pending drops: 0,
    Throttle drops: 0,
  V3 Input:
    Unknown security models: 0, Invalid messages: 0
    Unknown pdu handlers: 0, Unavailable contexts: 0
    Unknown contexts: 0, Unsupported security levels: 1
    Not in time windows: 0, Unknown user names: 0
    Unknown engine ids: 44, Wrong digests: 23, Decryption errors: 0
  Output:
    Packets: 246093, Too bigs: 0, No such names: 31561,
    Bad values: 0, General errors: 2,
    Get requests: 0, Get nexts: 0, Set requests: 0,
    Get responses: 246025, Traps: 0
```

show snmp v3

Syntax show snmp v3
<access <brief | detail> | community | general | groups | notify <filter> | target <address | parameters> | users>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display the Simple Network Management Protocol version 3 (SNMPv3) operating configuration.

Options none—Display all of the SNMPv3 operating configuration.

access—(Optional) Display SNMPv3 access information.

brief | detail—(Optional) Display brief or detailed information about SNMPv3 access information.

community—(Optional) Display SNMPv3 community information.

general—(Optional) Display SNMPv3 general information.

groups—(Optional) Display SNMPv3 security-to-group information.

notify <filter>—(Optional) Display SNMPv3 notify and, optionally, notify filter information.

target <address | parameters>—(Optional) Display SNMPv3 target and, optionally, either target address or target parameter information.

users—(Optional) Display SNMPv3 user information.

Additional Information To edit the default display of the **show snmp v3** command, specify options in the **show** statement at the **[edit snmp v3]** hierarchy level.

Required Privilege Level view

List of Sample Output **show snmp v3 on page 3400**

Output Fields Table 449 on page 3400 describes the output fields for the **show snmp v3** command. Output fields are listed in the approximate order in which they appear.

Table 449: show snmp v3 Output Fields

Field Name	Field Description
Access control	Information about access control: • **Group**—Group name for which the configured access privileges apply. The group, together with the context prefix and the security model and security level, forms the index for this table. • **Context prefix**—SNMPv3 context for which the configured access privileges apply. • **Security model/level**—Security model and security level for which the configuration access privileges apply. • **Read view**—Identifies the MIB view applied to SNMPv3 read operations. • **Write view**—Identifies the MIB view applied to SNMPv3 write operations. • **Notify view**—Identifies the MIB view applied to outbound SNMP notifications.
Engine	Information about local engine configuration: • **Local engine ID**—Identifier that uniquely and unambiguously identifies the local SNMPv3 engine. • **Engine boots**—Number of times the local SNMPv3 engine has rebooted or reinitialized since the engine ID was last changed. • **Engine time**—Number of seconds since the local SNMPv3 engine was last rebooted or reinitialized. • **Max msg size**—Maximum message size the sender can accommodate.
Engine ID	Information about engine ID: • **Local engine ID**—Identifier that uniquely and unambiguously identifies the local SNMPv3 engine. • **Engine boots**—Number of times the local SNMPv3 engine has rebooted or reinitialized since the engine ID was last changed. • **Engine time**—Number of seconds since the local SNMPv3 engine was last rebooted or reinitialized. • **Max msg size**—Maximum message size the sender can accommodate. • **Engine ID**—SNMPv3 engine ID associated with each user. • **User**—SNMPv3 user. • **Auth/Priv**—Authentication and encryption algorithm available for use by each user. • **Storage**—Indicates whether a user is saved to the configuration file (nonvolatile) or not (volatile). Applies only to users with active status. • **Status**—Status of the conceptual row. Only rows with an active status are used by the SNMPv3 engine.
Group name	Name of the group to which this entry belongs.
Security model	Identifies the security model context for the security name.
Security name	Used with the security model; identifies a specific security name instance. Each security model/security name combination can be assigned to a specific group.
Storage type	Indicates whether a user is saved to the configuration file (nonvolatile) or not (volatile). Applies only to users with active status.
Status	Status of the conceptual row. Only rows with active status are used by the SNMPv3 engine.

show snmp v3 `user@host> show snmp v3`

```
Local engine ID: 80 00 0a 4c e04 31 32 33 34
Engine boots:          38
Engine time:        64583 seconds
Max msg size:        2048 bytes

Engine ID: local
    User                           Auth/Priv   Storage      Status
    user1                           md5/des    nonvolatile  active
    user2                           sha/none   nonvolatile  active
    user3                          none/none   nonvolatile  active

Engine ID: 81 00 0a 4c 04 64 64 64 64
    User                           Auth/Priv   Storage      Status
    UNEW                            md5/none   nonvolatile  active
Group name           Security  Security              Storage      Status
                     model     name                  type
g1                   usm       user1                 nonvolatile  active
g2                   usm       user2                 nonvolatile  active
g3                   usm       user3                 nonvolatile  active

Access control:
Group                Context Security      Read       Write     Notify
                     prefix  model/level   view       view      view
g1                           usm/privacy   v1         v1
g2                           usm/authent   v1         v1
g3                           usm/none      v1         v1
```

CHAPTER 129

Real-Time Performance Monitoring (RPM)

RPM—Overview

Understanding Real-Time Performance Monitoring on J-EX Series Switches

Real-time performance monitoring (RPM) enables you to configure active probes to track and monitor traffic across the network and investigate network problems. You can use RPM with J-EX Series Switches.

The ways in which you can use RPM include:

- Monitor time delays between devices.
- Monitor time delays at the protocol level.
- Set thresholds to trigger SNMP traps when values are exceeded.

  You can configure thresholds for round-trip time, ingress or egress delay, standard deviation, jitter, successive lost probes, and total lost probes per test. (SNMP trap results are stored in **pingResultsTable**, **jnxPingResultsTable**, **jnxPingProbeHistoryTable**, and **pingProbeHistoryTable**.)
- Determine automatically whether a path exists between a host router or switch and its configured BGP neighbors. You can view the results of the discovery using an SNMP client.
- Use the history of the most recent 50 probes to analyze trends in your network and predict future needs.

RPM provides MIB support with extensions for RFC 2925, *Definitions of Managed Objects for Remote Ping, Traceroute, and Lookup Operations*.

This topic includes:

RPM Packet Collection

Probes collect packets per destination and per application, including PING Internet Control Message Protocol (ICMP) packets, User Datagram Protocol and Transmission Control Protocol (UDP/TCP) packets with user-configured ports, user-configured Differentiated Services code point (DSCP) type-of-service (ToS) packets, and Hypertext Transfer Protocol (HTTP) packets.

Tests and Probe Types

A test can contain multiple probes. The probe type specifies the packet and protocol contents of the probe.

J-EX Series switches support the following tests and probe types:

- Ping tests:
  - ICMP echo probe

    - ICMP timestamp probe
- HTTP tests:
    - HTTP get probe (not available for BGP RPM services)
    - HTTP get metadata probe
- UDP and TCP tests with user-configured ports:
    - UDP echo probe
    - TCP connection probe
    - UDP timestamp probe

Hardware Timestamps

To account for latency in the communication of probe messages, you can enable timestamping of the probe packets (hardware timestamps). If hardware timestamps are not configured, then timers are generated at the software level and are less accurate than they would have been with hardware timestamps.

NOTE: J-EX Series switches support hardware timestamps for UDP and ICMP probes. J-EX Series switches do not support hardware timestamps for HTTP or TCP probes.

You can timestamp the following RPM probes to improve the measurement of latency or jitter:

- ICMP ping
- ICMP ping timestamp
- UDP ping
- UDP ping timestamp

You should configure the requester with hardware timestamps (see Figure 87 on page 3406) to get more meaningful results than you would get without the timestamps. The responder does not need to be configured to support hardware timestamps. If the responder supports hardware timestamps, it will timestamp the RPM probes. If the responder does not support hardware timestamps, RPM can only report round-trip measurements that include the processing time on the responder.

NOTE: Hardware timestamps are supported on all J-EX Series switches.

Figure 87 on page 3406 shows the timestamps:

Figure 87: RPM Timestamps

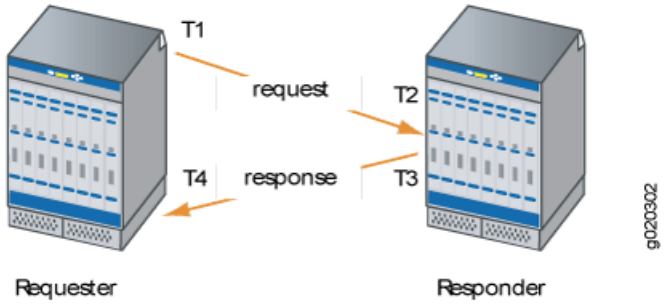

- T1 is the time the packet leaves the requester port.
- T2 is the time the responder receives the packet.
- T3 is the time the responder sends the response.
- T4 is the time the requester receives the response.

The round-trip time is (T2 – T1) + (T4 – T3). If the responder does not support hardware timestamps, then the round-trip time is (T4 – T1) / 2, and thus includes the processing time of the responder.

You can use RPM probes to find the following time measurements:

- Minimum round-trip time
- Maximum round-trip time
- Average round-trip time
- Standard deviation of the round-trip time
- Jitter of the round-trip time—Difference between the minimum and maximum round-trip time

NOTE: Configure hardware timestamps on the requester by including the hardware-timestamp statement at the [edit services rpm probe *probe-owner* test *test-name*] hierarchy level.

The RPM feature provides a configuration option to set one-way hardware timestamps. Use one-way timestamps when you want information about one-way time, rather than round-trip times, for packets to traverse the network between the requester and the responder. As shown in Figure 87 on page 3406, one-way timestamps represent the time T2 – T1 and the time from T4 – T3. Use one-way timestamps when you want to gather information about delay in each direction and to find egress and ingress jitter values.

NOTE: For correct one-way measurement, the clocks of the requester and responder must be synchronized. If the clocks are not synchronized, one-way jitter measurements and calculations can include significant variations, in some cases orders of magnitude greater than the round-trip times.

When you enable one-way timestamps in a probe, the following one-way measurements are reported:

- Minimum, maximum, standard deviation, and jitter measurements for egress and ingress times
- Number of probes sent
- Number of probe responses received
- Percentage of lost probes

Limitations of RPM

- Two-way Active Measurement Protocol (TWAMP) is not supported on J-EX Series switches.
- J-EX Series switches do not support user-configured class-of-service (CoS) classifiers or prioritization of RPM packets over regular data packets received on an input interface.
- Timestamps:
  - If the responder does not support hardware timestamps, RPM can only report the round-trip measurements and cannot calculate round-trip jitter.
  - J-EX Series switches do not support hardware timestamps for HTTP and TCP probes.
  - Timestamps apply only to IPv4 traffic.

Related Documentation

- For further details about RPM including configuration procedures, see *Junos OS Services Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/
- Configuring Real-Time Performance Monitoring (J-Web Procedure) on page 3407
- Configuring SNMP (J-Web Procedure) on page 3309
- Monitoring Network Traffic Using Traceroute on page 3515

Configuring Real-Time Performance Monitoring (RPM)

- Configuring Real-Time Performance Monitoring (J-Web Procedure) on page 3407
- Configuring the Interface for RPM Timestamping for Client/Server on a J-EX Series Switch (CLI Procedure) on page 3414

Configuring Real-Time Performance Monitoring (J-Web Procedure)

Real-time performance monitoring (RPM) in J-EX Series switches enables you to configure and send probes to a specified target and monitor the analyzed results to determine packet loss, round-trip time, and jitter. Jitter is the difference in relative transit time between two consecutive probes. You can set up probe owners and configure one or more performance probe tests under each probe owner.

The ways in which you can use RPM include:

- Monitor time delays between devices.
- Monitor time delays at the protocol level.
- Set thresholds to trigger SNMP traps when threshold values are exceeded. You can configure thresholds for round-trip time, ingress or egress delay, standard deviation, jitter, successive lost probes, and total lost probes per test.
- Determine automatically whether a path exists between a host switch and its configured Border Gateway Protocol (BGP) neighbors. You can view the results of the discovery using an SNMP client.
- Use the history of the most recent 50 probes to analyze trends in your network and predict future needs.

Probes collect packets per destination and per application, including PING Internet Control Message Protocol (ICMP) packets, User Datagram Protocol and Transmission Control Protocol (UDP/TCP) packets with user-configured ports, user-configured Differentiated Services code point (DSCP) type-of-service (ToS) packets, and Hypertext Transfer Protocol (HTTP) packets.

J-EX Series switches support the following tests and probe types:

- Ping tests:
  - ICMP echo
  - ICMP timestamp
- HTTP tests:
  - HTTP get (not available for BGP RPM services)
- UDP and TCP tests with user-configured ports:
  - UDP echo
  - TCP connection
  - UDP timestamp

To account for latency in the communication of probe messages, you can enable timestamping of the probe packets. You should configure both the requester and the responder to timestamp the RPM packets. The RPM features provides an additional configuration option to set one-way hardware timestamps. Use one-way timestamps when you want information about one-way, rather than round-trip, times for packets to traverse the network between the requester and the responder.

NOTE:

- J-EX Series switches support hardware timestamps for UDP and ICMP probes. J-EX Series switches do not support hardware timestamps for HTTP or TCP probes.
- If the responder does not support hardware timestamps, RPM can only report the round-trip measurements, it cannot calculate round-trip jitter.
- In J-EX Series switches timestamps apply only to IPv4 traffic.

To configure RPM using the J-Web interface:

1. Select **Troubleshoot** > **RPM** > **Configure RPM** .
2. In the **Configure RPM** page, enter information as specified in Table 450 on page 3409.
   a. Click **Add** to set up the **Owner Name** and **Performance Probe Tests**.
   b. Select a probe owner from **Probe Owners** list and click **Delete** to remove the selected probe owner
   c. Double-click one of the probe owners in **Probe Owners** list to display the list of performance probe tests.
   d. Double-click one of the performance probe tests to edit the test parameters.
3. Enter the **Maximum Number of Concurrent Probes** and specify the **Probe Servers**.
4. Click **Apply** to apply the RPM probe settings.

Table 450: RPM Probe Owner, Concurrent Probes, and Probe Servers Configuration Fields

Field	Function	Your Action
Probe Owners	Identifies a owner for whom one or more RPM tests are configured. In most implementations, the owner name identifies a network on which a set of tests is being run.	1. Click **Add** and type an owner name. 2. In **Performance Probe Tests**, click **Add** to define the RPM test parameters. See Table 451 on page 3410 for information on configuring RPM test parameters. 3. Click **OK** to save the settings or **Cancel** to exit from the window without saving the changes.
Maximum Number of Concurrent Probes	Specifies the maximum number of concurrent probes allowed.	Type a number from 1 through 500.

Table 450: RPM Probe Owner, Concurrent Probes, and Probe Servers Configuration Fields *(continued)*

Field	Function	Your Action
Probe Servers	Specifies the servers that act as receivers and transmitters for the probes.	Set up the following servers: • TCP Probe Server—Specifies the port on which the device is to receive and transmit TCP probes. Type the number 7 (a standard TCP port number) or a port number from 49160 through 65535. • UDP Probe Server—Specifies the port on which the device is to receive and transmit UDP probes. Type the number 7 (a standard TCP port number) or a port number from 49160 through 65535.

Table 451: Performance Probe Tests Configuration Fields

Field	Function	Your Action
Identification		
Test Name	Identifies the RPM test.	Type a test name.
Target (Address or URL)	Specifies the IP address or the URL of the probe target.	Type the IP address in dotted decimal notation or the URL of the probe target. If the target is a URL, type a fully formed URL that includes **http://**.
Source Address	Specifies the IP address to be used as the probe source address.	Type the source address to be used for the probe. If you do not supply this value, the packet uses the outgoing interface's address as the probe source address.
Routing Instance	Specifies the routing instance over which the probe is sent.	Type the routing instance name. The routing instance applies only to **icmp-ping** and **icmp-ping-timestamp** probe types. The default routing instance is **inet.0**.
History Size	Specifies the number of probe results to be saved in the probe history.	Type a number from 0 through 255. The default history size is 50.
Request Information		

Table 451: Performance Probe Tests Configuration Fields *(continued)*

Field	Function	Your Action
Probe Type	Specifies the type of probe to send as part of the test.	Select a probe type from the list: • **http-get** • **http-get-metadata** • **icmp-ping** • **icmp-ping-timestamp** • **tcp-ping** • **udp-ping** • **udp-ping-timestamp**
Interval	Sets the wait time (in seconds) between probe transmissions.	Type a number from 1 through 255 .
Test Interval	Sets the wait time (in seconds) between tests.	Type a number from 0 through 86400 .
Probe Count	Sets the total number of probes to be sent for each test.	Type a number from 1 through 15.
Moving Average Size	Specifies the number of samples to be used in the statistical calculation operations to be performed across a number of the most recent samples.	Type a number from 0 through 255.
Destination Port	Specifies the TCP or UDP port to which probes are sent. To use TCP or UDP probes, you must configure the remote server as a probe receiver. Both the probe server and the remote server must be Dell PowerConnect network devices configured to receive and transmit RPM probes on the same TCP or UDP port.	Type the number 7 (a standard TCP or UDP port number) or a port number from 49160 through 65535.
DSCP Bits	Specifies the Differentiated Services code point (DSCP) bits. This value must be a valid 6-bit pattern.	Type a valid 6-bit pattern.
Data Size	Specifies the size (in bytes) of the data portion of the ICMP probes.	Type a number from 0 through 65507.
Data Fill	Specifies the hexadecimal value of the data portion of the ICMP probes.	Type a hexadecimal value from 1h through 800h .
Hardware Timestamp		
One Way Hardware Timestamp	Enables one-way hardware timestamp.	To enable timestamping, select the check box.

Table 451: Performance Probe Tests Configuration Fields *(continued)*

Field	Function	Your Action
Destination Interface	Enables hardware timestamp on the specified interface.	Select an interface from the list.
Maximum Probe Thresholds		
Successive Lost Probes	Sets the number of probes that can be lost successively, if exceeded, triggers a probe failure and generates a system log message.	Type a number from 0 through 15.
Lost Probes	Sets the number of probes that can be lost , if exceeded, triggers a probe failure and generates a system log message.	Type a number from 0 through 15.
Round Trip Time	Sets the round-trip time (in microseconds), from the switch to the remote server, if exceeded, triggers a probe failure and generates a system log message.	Type a number from 0 through 60000000.
Jitter	Sets the jitter (in microseconds), if exceeded, triggers a probe failure and generates a system log message.	Type a number from 0 through 60000000.
Standard Deviation	Sets the maximum allowable standard deviation (in microseconds), if exceeded, triggers a probe failure and generates a system log message.	Type a number from 0 through 60000000.
Egress Time	Sets the one-way time (in microseconds), from the switch to the remote server, if exceeded, triggers a probe failure and generates a system log message.	Type a number from 0 through 60000000.
Ingress Time	Sets the one-way time (in microseconds), from the remote server to the switch, if exceeded, triggers a probe failure and generates a system log message.	Type a number from 0 through 60000000 (microseconds).
Jitter Egress Time	Sets the outbound-time jitter (in microseconds), if exceeded triggers a probe failure and generates a system log message.	Type a number from 0 through 60000000.
Jitter Ingress Time	Sets the inbound-time jitter (in microseconds), if exceeded, triggers a probe failure and generates a system log message.	Type a number from 0 and 60000000.

Table 451: Performance Probe Tests Configuration Fields *(continued)*

Field	Function	Your Action
Egress Standard Deviation	Sets the maximum allowable standard deviation of outbound times (in microseconds), if exceeded, triggers a probe failure and generates a system log message.	Type a number from 0 through 60000000.
Ingress Standard Deviation	Sets the maximum allowable standard deviation of inbound times (in microseconds), if exceeded, triggers a probe failure and generates a system log message.	Type a number from 0 through 60000000.
Traps		
Egress Jitter Exceeded	Generates SNMP traps when the threshold for jitter in outbound time is exceeded.	• To enable SNMP traps for this condition, select the check box. • To disable SNMP traps, clear the check box.
Egress Standard Deviation Exceeded	Generates SNMP traps when the threshold for standard deviation in outbound times is exceeded.	• To enable SNMP traps for this condition, select the check box. • To disable SNMP traps, clear the check box.
Egress Time Exceeded	Generates SNMP traps when the threshold for maximum outbound time is exceeded.	• To enable SNMP traps for this condition, select the check box. • To disable SNMP traps, clear the check box.
Ingress Jitter Exceeded	Generates SNMP traps when the threshold for jitter in inbound time is exceeded.	• To enable SNMP traps for this condition, select the check box. • To disable SNMP traps, clear the check box.
Ingress Standard Deviation Exceeded	Generates SNMP traps when the threshold for standard deviation in inbound times is exceeded.	• To enable SNMP traps for this condition, select the check box. • To disable SNMP traps, clear the check box.
Ingress Time Exceeded	Generates SNMP traps when the threshold for maximum inbound time is exceeded.	• To enable SNMP traps for this condition, select the check box. • To disable SNMP traps, clear the check box.
Jitter Exceeded	Generates SNMP traps when the threshold for jitter in round-trip time is exceeded.	• To enable SNMP traps for this condition, select the check box. • To disable SNMP traps, clear the check box.

Table 451: Performance Probe Tests Configuration Fields *(continued)*

Field	Function	Your Action
Probe Failure	Generates SNMP traps when the threshold for the number of successive lost probes is exceeded.	• To enable SNMP traps for this condition, select the check box. • To disable SNMP traps, clear the check box.
RTT Exceeded	Generates SNMP traps when the threshold for maximum round-trip time is exceeded.	• To enable SNMP traps for this condition, select the check box. • To disable SNMP traps, clear the check box.
Standard Deviation Exceeded	Generates SNMP traps when the threshold for standard deviation in round-trip times is exceeded.	• To enable SNMP traps for this condition, select the check box. • To disable SNMP traps, clear the check box.
Test Completion	Generates SNMP traps when a test is completed.	• To enable SNMP traps for this condition, select the check box. • To disable SNMP traps, clear the check box.
Test Failure	Generates SNMP traps when the threshold for the total number of lost probes is exceeded.	• To enable SNMP traps for this condition, select the check box. • To disable SNMP traps, clear the check box.

Related Documentation

- Configuring SNMP (J-Web Procedure) on page 3309
- Viewing Real-Time Performance Monitoring Information on page 3416

Configuring the Interface for RPM Timestamping for Client/Server on a J-EX Series Switch (CLI Procedure)

Use real-time performance monitoring (RPM) to configure active probes to track and monitor traffic across the network and to investigate network problems. To configure basic RPM probes on the J-EX Series switch, you must configure the probe owner, the test, and the specific parameters of the RPM probe.

You can also set a timestamp to improve the measurement of latency or jitter. The probe is timestamped by the device originating the probe (the RPM client). If you do not enable hardware timestamps, the timer values are set. You should configure both the RPM client (the requester) and the RPM server (the responder) to timestamp the RPM packets. However, if the RPM server does not support hardware timestamps, RPM can only report the round-trip measurements.

Timestamps apply only to IPv4 traffic.

You can enable hardware timestamps for the following RPM probe types:

- icmp-ping
- icmp-ping-timestamp
- udp-ping
- udp-ping-timestamp

To configure RPM probes and enable hardware timestamping:

1. Specify the probe owner:

   ```
   [edit services rpm]
   user@switch# set probe owner
   ```

2. Specify a test name. A test represents the range of probes over which the standard deviation, average, and jitter are calculated.

   ```
   [edit services rpm probe owner]
   user@switch# set test test-name
   ```

3. Specify the packet and protocol contents of the probe:

   ```
   [edit services rpm probe owner test test-name]
   user@switch# set probe-type type
   ```

4. Specify the destination IPv4 address to be used for the probes:

   ```
   [edit services rpm probe owner test test-name]
   user@switch# set target address address
   ```

5. Specify the number of probes within a test:

   ```
   [edit services rpm probe owner test test-name]
   user@switch# set probe-count count
   ```

6. Specify the time, in seconds, to wait between sending packets:

   ```
   [edit services rpm probe owner test test-name]
   user@switch# set probe-interval interval
   ```

7. Specify the time, in seconds, to wait between tests:

   ```
   [edit services rpm probe owner test test-name]
   user@switch# set test-interval interval
   ```

8. Specify the source IP address to be used for probes. If the source IP address is not one of the switch's assigned addresses, the packet uses the outgoing interface's address as its source.

   ```
   [edit services rpm probe owner test test-name]
   user@switch# set source-address address
   ```

9. Specify the value of the Differentiated Services (DiffServ) field within the IP header. The DiffServ code point (DSCP) bits value must be set to a valid 6-bit pattern.

   ```
   [edit services rpm probe owner test test-name]
   user@switch# set dscp-code-point dscp-bits
   ```

10. If you are using ICMP probes, specify the size of the data portion of ICMP probes:

    ```
    [edit services rpm probe owner test test-name]
    ```

```
user@switch# set data-size size
```

11. Enable hardware timestamping of RPM probe messages:

```
[edit services rpm probe owner test test-name]
user@switch# set hardware-timestamp
```

Related Documentation

- Configuring Real-Time Performance Monitoring (J-Web Procedure) on page 3407
- Understanding Real-Time Performance Monitoring on J-EX Series Switches on page 3404
- For details on these configuration statements, see the *Junos OS Services Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/index.html.

Verifying Real-Time Performance Monitoring

- Viewing Real-Time Performance Monitoring Information on page 3416

Viewing Real-Time Performance Monitoring Information

Real-time performance monitoring (RPM) on J-EX Series switches enables you to configure and send probes to a specified target and monitor the analyzed results to determine packet loss, round-trip time, and jitter. The J-Web interface provides a graphical view of RPM information for J-EX Series switches.

To view the RPM information using the J-Web interface:

1. Select **Troubleshoot**>**RPM**>**View RPM**.
2. Select the **Round Trip Time** check box to display the graph with round-trip time included. Clear the check-box to view the graph without the round-trip time.
3. From the **Refresh Time** list, select a refresh time interval for the graph.

Related Documentation

- Configuring Real-Time Performance Monitoring (J-Web Procedure) on page 3407

Operational Mode Commands for Real-Time Performance Monitoring

show services rpm active-servers

Syntax show services rpm active-servers

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display the protocols and corresponding ports for which a router or switch is configured as a real-time performance monitoring (RPM) server.

Options This command has no options.

Required Privilege Level view

List of Sample Output **show services rpm active-servers on page 3417**

Output Fields Table 452 on page 3417 lists the output fields for the **show services rpm active-servers** command. Output fields are listed in the approximate order in which they appear.

Table 452: show services rpm active-servers Output Fields

Field Name	Field Description
Protocol	Protocol configured on the receiving probe server. The protocol can be the User Datagram Protocol (UDP) or the Transmission Control Protocol (TCP).
Port	Port configured on the receiving probe server.
Destination interface name	Output interface name for the probes.

show services rpm active-servers

```
user@host> show services rpm active-servers
    Protocol: TCP, Port: 50000, Destination interface name: lt-0/0/0.0
    Protocol: UDP, Port: 50001, Destination interface name: lt-0/0/0.0
```

show services rpm history-results

Syntax show services rpm history-results
<brief | detail>
<owner *owner*>
<since *time*>
<test *name*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display standard information about the results of the last 50 probes for each real-time performance monitoring (RPM) instance.

Options none—Display the results of the last 50 probes for all RPM instances.

brief | detail—(Optional) Display the specified level of output.

owner *owner*—(Optional) Display information for the specified probe owner.

since *time*—(Optional) Display information from the specified time. Specify time as ***yyyy-mm-dd.hh*:*mm*:*ss***.

test *name*—(Optional) Display information for the specified test.

Required Privilege Level view

List of Sample Output **show services rpm history-results on page 3419**
show services rpm history-results detail on page 3420

Output Fields Table 453 on page 3418 lists the output fields for the **show services rpm history-results** command. Output fields are listed in the approximate order in which they appear.

Table 453: show services rpm history-results Output Fields

Field Name	Field Description	Level of Output
Owner	Probe owner.	All levels
Test	Name of a test for a probe instance.	All levels
Probe received	Timestamp when the probe result was determined.	All levels
Round trip time	Average ping round-trip time (RTT), in microseconds.	All levels
Probe results	Result of a particular probe performed by a remote host. The following information is contained in the results: • **Response received**—Timestamp when the probe result was determined. • **Rtt**—Average ping round-trip time (RTT), in microseconds.	**detail**

Table 453: show services rpm history-results Output Fields *(continued)*

Field Name	Field Description	Level of Output
Results over current test	Displays the results for the current test by probe at the time each probe was completed, as well as the status of the current test at the time the probe was completed.	**detail**
Probes sent	Number of probes sent with the current test.	**detail**
Probes received	Number of probe responses received within the current test.	**detail**
Loss percentage	Percentage of lost probes for the current test.	**detail**
Measurement	Increment of measurement. Possible values are round-trip time delay and, for the probe type icmp-pin-timestamp, the egress and ingress delay: • **Minimum**—Minimum RTT, ingress delay, or egress delay measured over the course of the current test. • **Maximum**—Maximum RTT, ingress delay, or egress delay measured over the course of the current test. • **Average**—Average RTT, ingress delay, or egress delay measured over the course of the current test. • **Jitter**—Difference, in microseconds, between the maximum and minimum RTT measured over the course of the current test. • **Stddev**—Standard deviation of the round-trip time, in microseconds, measured over the course of the current test.	**detail**

show services rpm history-results

```
user@host> show services rpm history-results
     Owner, Test                 Probe received              Round trip time
     flintstone, 0               Tue Dec 28 15:56:22 2004          158 usec
     flintstone, 0               Tue Dec 28 15:56:23 2004          218 usec
     flintstone, 0               Tue Dec 28 15:56:24 2004          161 usec
     flintstone, 0               Tue Dec 28 15:56:25 2004          184 usec
     flintstone, 0               Tue Dec 28 15:56:30 2004          332 usec
     flintstone, 0               Tue Dec 28 15:56:31 2004          132 usec
     flintstone, 0               Tue Dec 28 15:56:32 2004          226 usec
     flintstone, 0               Tue Dec 28 15:56:33 2004          191 usec
     flintstone, 0               Tue Dec 28 15:56:34 2004          179 usec
     flintstone, 0               Tue Dec 28 15:56:39 2004          217 usec
     flintstone, 0               Tue Dec 28 15:56:40 2004          141 usec
     flintstone, 0               Tue Dec 28 15:56:41 2004          230 usec
     flintstone, 0               Tue Dec 28 15:56:42 2004          248 usec
     flintstone, 0               Tue Dec 28 15:56:43 2004          234 usec
     flintstone, 0               Tue Dec 28 15:56:48 2004          251 usec
     flintstone, 0               Tue Dec 28 15:56:49 2004          134 usec
     flintstone, 0               Tue Dec 28 15:56:50 2004          272 usec
     flintstone, 0               Tue Dec 28 15:56:51 2004          181 usec
     flintstone, 0               Tue Dec 28 15:56:52 2004          216 usec
     flintstone, 0               Tue Dec 28 15:56:57 2004          227 usec
     flintstone, 0               Tue Dec 28 15:56:58 2004          133 usec
```

show services rpm history-results detail

```
user@host> show services rpm history-results detail
Owner: flintstone, Test: 0
     Probe results:
       Response received, Tue Dec 28 15:56:39 2004
       Rtt: 217 usec
     Results over current test:
       Probes sent: 1, Probes received: 1, Loss percentage: 0
       Measurement: Round trip time
         Minimum: 217 usec, Maximum: 217 usec, Average: 217 usec,
         Jitter: 0 usec, Stddev: 0 usec

   Owner: flintstone, Test: 0
     Probe results:
       Response received, Tue Dec 28 15:56:40 2004
       Rtt: 141 usec
     Results over current test:
       Probes sent: 2, Probes received: 2, Loss percentage: 0
       Measurement: Round trip time
         Minimum: 141 usec, Maximum: 217 usec, Average: 179 usec,
         Jitter: 76 usec, Stddev: 38 usec

   Owner: flintstone, Test: 0
     Probe results:
       Response received, Tue Dec 28 15:56:41 2004
       Rtt: 230 usec
     Results over current test:
       Probes sent: 3, Probes received: 3, Loss percentage: 0
       Measurement: Round trip time
         Minimum: 141 usec, Maximum: 230 usec, Average: 196 usec,
         Jitter: 89 usec, Stddev: 39 usec

   Owner: flintstone, Test: 0
     Probe results:
       Response received, Tue Dec 28 15:56:42 2004
       Rtt: 248 usec
     Results over current test:
       Probes sent: 4, Probes received: 4, Loss percentage: 0
       Measurement: Round trip time
         Minimum: 141 usec, Maximum: 248 usec, Average: 209 usec,
         Jitter: 107 usec, Stddev: 41 usec
```

show services rpm probe-results

Syntax
show services rpm probe-results
<owner *owner*>
<test *name*>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Display the results of the most recent real-time performance monitoring (RPM) probes.

Options
none—Display all results of the most recent RPM probes.

owner *owner*—(Optional) Display information for the specified probe owner.

test *name*—(Optional) Display information for the specified test.

Required Privilege Level
view

List of Sample Output
show services rpm probe-results on page 3424
show services rpm probe-results (BGP Neighbor Discovery) on page 3425

Output Fields
Table 454 on page 3421 lists the output fields for the **show services rpm probe-results** command. Output fields are listed in the approximate order in which they appear.

Table 454: show services rpm probe-results Output Fields

Field Name	Field Description
Owner	Owner name. When you configure the probe owner statement at the **[edit services rpm]** hierarchy level, this field displays the configured owner name. When you configure BGP neighbor discovery through RPM, the output for this field is **Rpm-Bgp-Owner.**
Test	Name of a test representing a collection of probes. When you configure the test test-name statement at the **[edit services rpm probe owner]** hierarchy level, the field displays the configured test name. When you configure BGP neighbor discovery through RPM, the output for this field is **Rpm-BGP-Test-***n*, where *n* is a cumulative number.
Target address	Destination address used for the probes.
Source address	Source address used for the probes.
Probe type	Protocol configured on the receiving probe server: **http-get, http-metadata-get, icmp-ping, icmp-ping-timestamp, tcp-ping, udp-ping,** or **udp-ping-timestamp**.
Test size	Number of probes within a test.

Table 454: show services rpm probe-results Output Fields *(continued)*

Field Name	Field Description
Routing Instance Name	(BGP neighbor discovery) Name of the configured (if any) routing instance, logical system name, or both, in which the probe is configured: • When a routing instance is defined within a logical system, the logical system name is followed by the routing instance name. A slash **(/)** is used to separate the two entities. For example, if the routing instance called **R1** is configured within the logical system called **LS**, the name in the output field is **LS/R1**. • When a routing instance is configured but the default logical system is used, the name in the output field is the name of the routing instance. • When a logical system is configured but the default routing instance is used, the name in the output field is the name of the logical system followed by **default**. A slash (**/**) is used to separate the two entities. For example, **LS/default**.
Probe results	Raw measurement of a particular probe sample done by a remote host. This data is provided separately from the calculated results. The following information is contained in the raw measurement: • **Response received**—Timestamp when the probe result was determined. • **Client and server hardware timestamps**—If timestamps are configured, an entry appears at this point. • **Rtt**—Average ping round-trip time (RTT), in microseconds. • **Egress jitter**—Egress jitter, in microseconds. • **Ingress jitter**—Ingress jitter, in microseconds. • **Round trip jitter**—Round-trip jitter, in microseconds. • **Egress interarrival jitter**—Egress interarrival jitter, in microseconds. • **Ingress interarrival jitter**—Ingress interarrival jitter, in microseconds. • **Round trip interarrival jitter**—Round-trip interarrival jitter, in microseconds.
Results over current test	Probes are grouped into tests, and the statistics are calculated for each test. If a test contains 10 probes, the average, minimum, and maximum results are calculated from the results of those 10 probes. If the command is issued while the test is in progress, the statistics use information from the completed probes. • **Probes sent**—Number of probes sent within the current test. • **Probes received**—Number of probe responses received within the current test. • **Loss percentage**—Percentage of lost probes for the current test. • **Measurement**—Measurement type. Possible values are round-trip time, positive round-trip jitter, negative round-trip jitter, egress time, positive egress jitter, negative egress jitter, ingress time, positive ingress jitter, negative ingress jitter, and, for the probe type **icmp-ping-timestamp**, the egress delay and ingress delay. For each measurement type, the following individual calculated results are provided: ◦ **Samples**—Number of probes. ◦ **Minimum**—Minimum RTT, ingress delay, or egress delay measured over the course of the current test. ◦ **Maximum**—Maximum RTT, ingress delay, or egress delay measured over the course of the current test. ◦ **Average**—Average RTT, ingress delay, or egress delay measured over the course of the current test. ◦ **Peak to peak**—Peak-to-peak difference, in microseconds. ◦ **Stddev**—Standard deviation, in microseconds. ◦ **Sum**—Statistical sum.

Table 454: show services rpm probe-results Output Fields *(continued)*

Field Name	Field Description
Results over last test	Results for the most recently completed test. If the command is issued while the first test is in progress, this information is not displayed • **Probes sent**—Number of probes sent for the most recently completed test. • **Probes received**—Number of probe responses received for the most recently completed test. • **Loss percentage**—Percentage of lost probes for the most recently completed test. • **Test completed**—Time the most recent test was completed. • **Measurement**—Measurement type. Possible values are round-trip time, positive round-trip jitter, negative round-trip jitter, egress time, positive egress jitter, negative egress jitter, ingress time, positive ingress jitter, negative ingress jitter, and, for the probe type **icmp-ping-timestamp**, the egress delay and ingress delay. For each measurement type, the following individual calculated results are provided: ◦ **Samples**—Number of probes. ◦ **Minimum**—Minimum RTT, ingress delay, or egress delay measured for the most recently completed test. ◦ **Maximum**—Maximum RTT, ingress delay, or egress delay measured for the most recently completed test. ◦ **Average**—Average RTT, ingress delay, or egress delay measured for the most recently completed test. ◦ **Peak to peak**—Peak-to-peak difference, in microseconds. ◦ **Stddev**—Standard deviation, in microseconds. ◦ **Sum**—Statistical sum.
Results over all tests	Displays statistics made for all the probes, independently of the grouping into tests, as well as statistics for the current test. • **Probes sent**—Number of probes sent in all tests. • **Probes received**—Number of probe responses received in all tests. • **Loss percentage**—Percentage of lost probes in all tests. • **Measurement**—Measurement type. Possible values are round-trip time, positive round-trip jitter, negative round-trip jitter, egress time, positive egress jitter, negative egress jitter, ingress time, positive ingress jitter, negative ingress jitter, and, for the probe types **icmp-ping-timestamp** and **udp-ping-timestamp**, the egress delay and ingress delay. For each measurement type, the following individual calculated results are provided: ◦ **Samples**—Number of probes. ◦ **Minimum**—Minimum RTT, ingress delay, or egress delay measured over the course of the current test. ◦ **Maximum**—Maximum RTT, ingress delay, or egress delay measured over the course of the current test. ◦ **Average**—Average RTT, ingress delay, or egress delay measured over the course of the current test. ◦ **Peak to peak**—Peak-to-peak difference, in microseconds. ◦ **Stddev**—Standard deviation, in microseconds. ◦ **Sum**—Statistical sum.

show services rpm probe-results

```
user@host> show services rpm probe-results
    Owner: ADSN-J4300.ADSN-J2300.D2, Test: 75300002
    Target address: 172.16.54.172, Source address: 10.206.0.1,
    Probe type: udp-ping-timestamp, Test size: 10 probes
    Probe results:
      Response received, Tue Feb  6 14:53:15 2007,
      Client and server hardware timestamps
      Rtt: 575 usec, Egress jitter: 5 usec, Ingress jitter: 8 usec,
      Round trip jitter: 12 usec, Egress interarrival jitter: 8 usec,
     Ingress interarrival jitter: 7 usec, Round trip interarrival jitter: 7 usec,

      Round trip interarrival jitter: 669 usec
    Results over current test:
      Probes sent: 10, Probes received: 10, Loss percentage: 0
      Measurement: Round trip time
        Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
        Peak to peak: 2054 usec, Stddev: 738 usec,  Sum: xxxx usec
      Measurement: Positive round trip jitter
        Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
        Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
      Measurement: Negative round trip jitter
        Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
        Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
      Measurement: Egress time
        Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
        Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
      Measurement: Positive Egress  jitter
        Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
        Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
      Measurement: Negative Egress  jitter
        Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
        Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
      Measurement: Ingress time
        Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
        Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
      Measurement: Positive Ingress  jitter
        Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
        Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
      Measurement: Negative Ingress  jitter
        Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
        Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
    Results over last test:
      Probes sent: 10, Probes received: 10, Loss percentage: 0
      Test completed on Tue Feb  6 14:53:16 2007
      Measurement: Round trip time
        Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
        Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
      Measurement: Positive round trip jitter
        Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
        Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
      Measurement: Negative round trip jitter
        Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
        Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
      Measurement: Egress time
        Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
        Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
      Measurement: Positive Egress  jitter
        Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
        Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
      Measurement: Negative Egress  jitter
        Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
```

```
    Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
  Measurement: Ingress time
    Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
    Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
  Measurement: Positive Ingress  jitter
    Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
    Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
  Measurement: Negative Ingress  jitter
    Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
    Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
Results over all tests:
  Probes sent: 560, Probes received: 560, Loss percentage: 0
  Measurement: Round trip time
    Samples: 560, Minimum: 805 usec, Maximum: 3114 usec, Average: 1756 usec,

    Peak to peak: 2309 usec, Stddev: 519 usec, Sum: xxxx usec
  Measurement: Positive round trip jitter
    Samples: 257, Minimum: 0 usec, Maximum: 2054 usec, Average: 597 usec,
    Peak to peak: 2054 usec, Stddev: 427 usec, Sum: xxxx usec
  Measurement: Negative round trip jitter
    Samples: 302, Minimum: 1 usec, Maximum: 1812 usec, Average: 511 usec,
    Peak to peak: 1811 usec, Stddev: 408 usec, Sum: xxxx usec
  Measurement: Egress time
    Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
    Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
  Measurement: Positive Egress  jitter
    Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
    Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
  Measurement: Negative Egress  jitter
    Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
    Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
  Measurement: Ingress time
    Samples: 10, Minimum: 805 usec, Maximum: 2859 usec, Average: 1644 usec,
    Peak to peak: 2054 usec, Stddev: 738 usec, Sum: xxxx usec
  Measurement: Positive Ingress  jitter
    Samples: 5, Minimum: 5 usec, Maximum: 2054 usec, Average: 876 usec,
    Peak to peak: 2049 usec, Stddev: 679 usec, Sum: xxxx usec
  Measurement: Negative Ingress  jitter
    Samples: 5, Minimum: 5 usec, Maximum: 1812 usec, Average: 926 usec,
    Peak to peak: 1807 usec, Stddev: 665 usec, Sum: xxxx usec
```

show services rpm probe-results (BGP Neighbor Discovery)

```
user@host> show services rpm probe-results
Owner: Rpm-Bgp-Owner, Test: Rpm-Bgp-Test-1
    Target address: 10.209.152.37, Probe type: icmp-ping, Test size: 5 probes
    Routing Instance Name: LS1/RI1
    Probe results:
      Response received, Fri Oct 28 05:20:23 2005
      Rtt: 662 usec
    Results over current test:
      Probes sent: 5, Probes received: 5, Loss percentage: 0
      Measurement: Round trip time
        Minimum: 529 usec, Maximum: 662 usec, Average: 585 usec,
        Jitter: 133 usec, Stddev: 53 usec
    Results over all tests:
      Probes sent: 5, Probes received: 5, Loss percentage: 0
      Measurement: Round trip time
        Minimum: 529 usec, Maximum: 662 usec, Average: 585 usec,
        Jitter: 133 usec, Stddev: 53 usec
```

CHAPTER 130

Ethernet OAM Link Fault Management

Ethernet OAM Link Fault Management—Overview

Understanding Ethernet OAM Link Fault Management for a J-EX Series Switch

The Junos OS for J-EX Series Switches allows the Ethernet interfaces on these switches to support the IEEE 802.3ah standard for the Operation, Administration, and Maintenance (OAM) of Ethernet in access networks. The standard defines OAM link fault management (LFM). You can configure IEEE 802.3ah OAM LFM on point-to-point Ethernet links that are connected either directly or through Ethernet repeaters. The IEEE 802.3ah standard meets the requirement for OAM capabilities even as Ethernet moves from being solely an enterprise technology to a WAN and access technology, and the standard remains backward-compatible with existing Ethernet technology.

Ethernet OAM provides the tools that network management software and network managers can use to determine how a network of Ethernet links is functioning. Ethernet OAM should:

- Rely only on the media access control (MAC) address or virtual LAN identifier for troubleshooting.
- Work independently of the actual Ethernet transport and function over physical Ethernet ports or a virtual service such as pseudowire.
- Isolate faults over a flat (or single operator) network architecture or nested or hierarchical (or multiprovider) networks.

The following OAM LFM features are supported on J-EX Series switches:

- Discovery and Link Monitoring

The discovery process is triggered automatically when OAM is enabled on the interface. The discovery process permits Ethernet interfaces to discover and monitor the peer on the link if it also supports the IEEE 802.3ah standard. You can specify the discovery mode used for IEEE 802.3ah OAM support. In active mode, the interface discovers and monitors the peer on the link if the peer also supports IEEE 802.3ah OAM functionality. In passive mode, the peer initiates the discovery process. After the discovery process has been initiated, both sides participate in discovery. The switch performs link monitoring by sending periodic OAM protocol data units (PDUs) to advertise OAM mode, configuration, and capabilities.

You can specify the number of OAM PDUs that an interface can miss before the link between peers is considered down.

- Remote Fault Detection

  Remote fault detection uses flags and events. Flags are used to convey the following: Link Fault means a loss of signal, Dying Gasp means an unrecoverable condition such as a power failure, and Critical Event means an unspecified vendor-specific critical event. You can specify the periodic OAM PDU sending interval for fault detection.The J-EX Series switch uses the Event Notification OAM PDU to notify the remote OAM device when a problem is detected. You can specify the action to be taken by the system when the configured link-fault event occurs.

- Remote Loopback Mode

  Remote loopback mode ensures link quality between the switch and a remote peer during installation or troubleshooting. In this mode, when the interface receives a frame that is not an OAM PDU or a pause frame, it sends it back on the same interface on which it was received. The link appears to be in the active state. You can use the returned loopback acknowledgement to test delay, jitter, and throughput.

  Junos OS can place a remote DTE into loopback mode (if remote loopback mode is supported by the remote DTE). When you place a remote DTE into loopback mode, the interface receives the remote loopback request and puts the interface into remote loopback mode. When the interface is in remote loopback mode, all frames except OAM PDUs are looped back without any changes made to the frames. OAM PDUs continue to be sent and processed.

Related Documentation

- Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 3431
- Example: Configuring Ethernet OAM Link Fault Management on J-EX Series Switches on page 3428

Example of Ethernet OAM Link Fault Management Configuration

- Example: Configuring Ethernet OAM Link Fault Management on J-EX Series Switches on page 3428

Example: Configuring Ethernet OAM Link Fault Management on J-EX Series Switches

The Junos OS for J-EX Series switches allows the Ethernet interfaces on these switches to support the IEEE 802.3ah standard for the Operation, Administration, and Maintenance

(OAM) of Ethernet in access networks. The standard defines OAM link fault management (LFM). You can configure IEEE 802.3ah OAM LFM on point-to-point Ethernet links that are connected either directly or through Ethernet repeaters.

This example describes how to enable and configure OAM LFM on a Gigabit Ethernet interface:

Requirements

This example uses the following hardware and software components:

- Two J-EX4200 switches connected directly

Overview and Topology

Junos OS for J-EX Series switches allows the Ethernet interfaces on these switches to support the IEEE 802.3ah standard for the Operation, Administration, and Maintenance (OAM) of Ethernet in access networks. The standard defines OAM link fault management (LFM). You can configure IEEE 802.3ah OAM LFM on point-to-point Ethernet links that are connected either directly or through Ethernet repeaters.

This example uses two J-EX4200 switches connected directly. Before you begin configuring Ethernet OAM LFM on two switches, connect the two switches directly through a trunk interface.

Configuring Ethernet OAM Link Fault Management on Switch 1

CLI Quick Configuration

To quickly configure Ethernet OAM LFM, copy the following commands and paste them into the switch terminal window:

```
[edit protocols oam ethernet link-fault-management]
set interface ge-0/0/0
set interface ge-0/0/0 link-discovery active
set interface ge-0/0/0 pdu-interval 800
set interface ge-0/0/0 remote-loopback
```

Step-by-Step Procedure

To configure Ethernet OAM LFM on switch 1:

1. Enable IEEE 802.3ah OAM support on an interface:

   ```
   [edit protocols oam ethernet link-fault-management]
   user@switch1# set interface ge-0/0/0
   ```

2. Specify that the interface initiates the discovery process by configuring the link discovery mode to **active**:

   ```
   [edit protocols oam ethernet link-fault-management]
   user@switch1# set interface ge-0/0/0 link-discovery active
   ```

3. Set the periodic OAM PDU-sending interval (in milliseconds) to 800 on switch 1:

```
[edit protocols oam ethernet link-fault-management]
user@switch1# set interface pdu-interval 800
```

4. Set a remote interface into loopback mode so that all frames except OAM PDUs are looped back without any changes made to the frames. Ensure that the remote DTE supports remote loopback mode. To set the remote DTE in loopback mode

```
[edit protocols oam ethernet link-fault-management]
user@switch1# set interface ge-0/0/0.0 remote-loopback
```

Results Check the results of the configuration:

```
[edit]
user@switch1# show
```

```
protocols {
    oam {
        ethernet {
            link-fault-management {
                interface ge-0/0/0 {
                    pdu-interval 800;
                    link-discovery active;
                    remote-loopback;
                }
            }
        }
    }
```

Configuring Ethernet OAM Link Fault Management on Switch 2

CLI Quick Configuration To quickly configure Ethernet OAM LFM on switch 2, copy the following commands and paste them into the switch terminal window:

```
[edit protocols oam ethernet link-fault-management ]
set interface ge-0/0/1
set interface ge-0/0/1 negotiation-options allow-remote-loopback
```

Step-by-Step Procedure To configure Ethernet OAM LFM on switch 2:

1. Enable OAM on the peer interface on switch 2:

```
[edit protocols oam ethernet link-fault-management]
user@switch2# set interface ge-0/0/1
```

2. Enable remote loopback support for the local interface:

```
[edit protocols oam ethernet link-fault-management]
user@switch2# set interface ge-0/0/1 negotiation-options allow-remote-loopback
```

Results Check the results of the configuration:

```
[edit]
user@switch2# show
```

```
protocols {
    oam {
        ethernet {
            link-fault-management {
```

```
            interface ge-0/0/1 {
              negotiation-options {
                allow-remote-loopback;
              }
            }
          }
        }
      }
```

Verification

Verifying That OAM LFM Has Been Configured Properly

Purpose Verify that OAM LFM has been configured properly.

Action Use the **show oam ethernet link-fault-management** command:

```
user@switch1#show oam ethernet link-fault-management
```

Sample Output

```
Interface: ge-0/0/0.0
Status: Running, Discovery state: Send Any
Peer address: 00:19:e2:50:3b:e1
Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
Remote entity information:
Remote MUX action: forwarding, Remote parser action: forwarding
Discovery mode: active, Unidirectional mode: unsupported
Remote loopback mode: supported, Link events: supported
Variable requests: unsupported
```

Meaning When the output displays the MAC address and the discover state is **Send Any**, it means that OAM LFM has been configured properly.

Related Documentation

- Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 3431
- Understanding Ethernet OAM Link Fault Management for a J-EX Series Switch on page 3427

Configuring Ethernet OAM Link Fault Management

Configuring Ethernet OAM Link Fault Management (CLI Procedure)

Ethernet OAM link fault management (LFM) can be used for physical link-level fault detection and management. The IEEE 802.3ah LFM works across point-to-point Ethernet links either directly or through repeaters.

To configure Ethernet OAM LFM using the CLI:

1. Enable IEEE 802.3ah OAM support on an interface:

   ```
   [edit protocols oam ethernet link-fault-management]
   user@switch# set interface interface-name
   ```

   NOTE: The remaining steps are optional. You can choose which of these features to configure for Ethernet OAM LFM on your switch.

2. Specify whether the interface or the peer initiates the discovery process by configuring the link discovery mode to **active** or **passive** (**active** = interface initiates; **passive** = peer initiates):

   ```
   [edit protocols oam ethernet link-fault-management]
   user@switch# set interface interface-name link-discovery active
   ```

3. Configure a periodic OAM PDU-sending interval (in milliseconds) for fault detection:

   ```
   [edit protocols oam ethernet link-fault-management]
   user@switch# set interface pdu-interval interval
   ```

4. Specify the number of OAM PDUs that an interface can miss before the link between peers is considered down:

   ```
   [edit protocols oam ethernet link-fault-management]
   user@switch# set interface interface-name pdu-threshold threshold-value
   ```

5. Configure event threshold values on an interface for the local errors that trigger the sending of link event TLVs:

   - Set the threshold value (in seconds) for sending frame-error events or taking the action specified in the action profile:

     ```
     [edit protocols oam ethernet link-fault-management]
     user@switch# set interface interface-name event-thresholds frame-error count
     ```

   - Set the threshold value (in seconds) for sending frame-period events or taking the action specified in the action profile:

     ```
     [edit protocols oam ethernet link-fault-management]
     user@switch#  set interface interface-name event-thresholds frame-period count
     ```

   - Set the threshold value (in seconds) for sending frame-period-summary events or taking the action specified in the action profile:

     ```
     [edit protocols oam ethernet link-fault-management]
     user@switch#  set interface interface-name event-thresholds frame-period-summary
     count
     ```

   - Set the threshold value (in seconds) for sending symbol-period events or taking the action specified in the action profile:

     ```
     [edit protocols oam ethernet link-fault-management]
     user@switch#  set interface interface-name event-thresholds symbol-period count
     ```

NOTE: You can disable the sending of link event TLVs.

To disable the sending of link event TLVs:

```
[edit protocols oam ethernet link-fault-management]
user@switch# set interface interface-name negotiation-options no-allow-link-events
```

6. Create an action profile to define event fault flags and thresholds to be taken when the link fault event occurs. Then apply the action profile to one or more interfaces. (You can also apply multiple action profiles to a single interface.)

   a. Name the action profile:

   ```
   [edit protocols oam ethernet link-fault-management]
   user@switch# set action-profile profile-name
   ```

   b. Specify actions to be taken by the system when the link fault event occurs:

   ```
   [edit protocols oam ethernet link-fault-management]
   user@switch# set action-profile profile-name  action syslog

   user@switch# set action-profile profile-name action link-down
   ```

   c. Specify events for the action profile:

   ```
   [edit protocols oam ethernet link-fault-management]
   user@switch# set action-profile profile-name event link-adjacency-loss
   ```

NOTE: For each action profile, you must specify at least one link event and one action. The actions are taken only when all of the events in the action profile are true. If more than one action is specified, all actions are executed. You can set a low threshold for a specific action such as logging the error and set a high threshold for another action such as system logging.

7. Set a remote interface into loopback mode so that all frames except OAM PDUs are looped back without any changes made to the frames. Set the remote DTE in loopback mode (the remote DTE must support remote-loopback mode) and then enable remote loopback support for the local interface.

   ```
   [edit protocols oam ethernet link-fault-management]
   user@switch# set interface interface-name remote-loopback

   user@switch# set interface interface-name negotiation-options allow-remote-loopback
   ```

Related Documentation

- Example: Configuring Ethernet OAM Link Fault Management on J-EX Series Switches on page 3428
- Understanding Ethernet OAM Link Fault Management for a J-EX Series Switch on page 3427

Configuration Statements for Ethernet OAM Link Fault Management

- [edit protocols] Configuration Statement Hierarchy on page 3434

[edit protocols] Configuration Statement Hierarchy

```
protocols {
  connections {
    remote-interface-switch connection-name {
      interface interface-name.unit-number;
      transmit-lsp label-switched-path;
      receive-lsp label-switched-path;
    }
  }
  dot1x {
    authenticator {
      authentication-profile-name profile-name;
      interface (all | [ interface-names ]) {
        disable;
        guest-vlan ( vlan-id | vlan-name);
        mac-radius <restrict>;
        maximum-requests number;
        no-reauthentication;
        quiet-period seconds;
        reauthentication {
          interval seconds;
        }
        retries number;
        server-fail (deny | permit | use-cache | vlan-id | vlan-name);
        server-reject-vlan (vlan-id | vlan-name);
        server-timeout seconds;
        supplicant (multiple | single | single-secure);
        supplicant-timeout seconds;
        transmit-period seconds;
      }
      static mac-address {
        interface interface-name;
        vlan-assignment (vlan-id |vlan-name);
      }
  }
  gvrp {
    <enable | disable>;
    interface (all | [interface-name]) {
      disable;
    }
    join-timer millseconds;
    leave-timer milliseconds;
    leaveall-timer milliseconds;
  }
  igmp-snooping {
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>
        <match regex>;
      flag flag (detail | disable | receive | send);
    }
```

```
    vlan (vlan-id | vlan-number) {
      data-forwarding {
        source {
          groups group-prefix;
        }
        receiver {
          source-vlans vlan-list;
          install ;
        }
      }
      disable {
        interface interface-name
      }
      immediate-leave;
      interface interface-name {
        group-limit limit;
        multicast-router-interface;
        static {
          group ip-address;
        }
      }
      proxy ;
      query-interval seconds;
      query-last-member-interval seconds;
      query-response-interval seconds;
      robust-count number;
    }
  }
  lldp {
    disable;
    advertisement-interval seconds;
    hold-multiplier number;
    interface (all | interface-name) {
      disable;
    }
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>
        <match regex>;
      flag flag (detail | disable | receive | send);
    }
  }
  lldp-med {
    disable;
    fast-start number;
    interface (all | interface-name) {
      disable;
      location {
        elin number;
        civic-based {
          what number;
          country-code code;
          ca-type {
            number {
              ca-value value;
            }
          }
```

```
        }
      }
    }
  }
  mpls {
    interface ( all | interface-name );
    label-switched-path lsp-name to remote-provider-edge-switch;
    path destination {
      <address | hostname> <strict | loose>
    }
  mstp {
    disable;
    bpdu-block-on-edge;
    bridge-priority priority;
    configuration-name name;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
      disable;
      bpdu-timeout-action {
        block;
        alarm;
      }
      cost  cost;
      edge;
      mode mode;
      no-root-port;
      priority priority;
    }
    max-age seconds;
    max-hops hops;
    msti msti-id {
      vlan (vlan-id | vlan-name);
      interface interface-name {
        disable;
        cost  cost;
        edge;
        mode mode;
        priority priority;
      }
    }
    revision-level revision-level;
    traceoptions {
      file filename <files number > <size size> <no-stamp | world-readable |
        no-world-readable>;
      flag flag;
    }
  }
  mvrp {
    disable
    interface (all | interface-name) {
      disable;
      join-timer milliseconds;
      leave-timer milliseconds;
      leaveall-timer milliseconds;
      registration (forbidden | normal);
```

```
    }
    no-dynamic-vlan;
    traceoptions {
      file filename <files number > <size size> <no-stamp | world-readable |
        no-world-readable>;
      flag flag;
    }
  }
  oam {
    ethernet{
      connectivity-fault-management {
        action-profile profile-name {
          default-actions {
            interface-down;
          }
        }
        linktrace {
          age (30m | 10m | 1m | 30s | 10s);
          path-database-size path-database-size;
        }
        maintenance-domain domain-name {
          level number;
          mip-half-function (none | default |explicit);
          name-format (character-string | none | dns | mac+2oct);
          maintenance-association ma-name {
            continuity-check {
              hold-interval minutes;
              interval (10m | 10s | 1m | 1s| 100ms);
              loss-threshold number;
            }
            mep mep-id {
              auto-discovery;
              direction down;
              interface interface-name;
              remote-mep mep-id {
                action-profile profile-name;
              }
            }
          }
        }
      }
      link-fault-management {
        action-profile profile-name;
          action {
            syslog;
            link-down;
          }
          event {
            link-adjacency-loss;
            link-event-rate;
            frame-error count;
            frame-period count;
            frame-period-summary count;
            symbol-period count;
          }
          interface interface-name {
```

```
            link-discovery (active | passive);
            pdu-interval interval;
            event-thresholds threshold-value;
            remote-loopback;
            event-thresholds {
              frame-errorcount;
              frame-period count;
              frame-period-summary count;
              symbol-period count;
            }
          }
          negotiation-options {
            allow-remote-loopback;
            no-allow-link-events;
          }
        }
      }
    }
    rstp {
      disable;
      bpdu-block-on-edge;
      bridge-priority priority;
      forward-delay seconds;
      hello-time seconds;
      interface (all | interface-name) {
        disable;
        bpdu-timeout-action {
          block;
          alarm;
        }
        cost  cost;
        edge;
        mode mode;
        no-root-port;
        priority priority;
      }
      max-age seconds;
      }
      traceoptions {
        file filename <files number > <size size> <no-stamp | world-readable |
          no-world-readable>;
        flag flag;
      }
    }
    sflow {
      agent-id
      collector {
        ip-address;
        udp-port port-number;
      }
      disable;
      interfaces interface-name {
        disable;
        polling-interval seconds;
        sample-rate number;
      }
```

```
    polling-interval seconds;
    sample-rate number;
    source-ip
  }
  stp {
    disable;
    bridge-priority priority;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
      disable;
      bpdu-timeout-action {
        block;
        alarm;
      }
      cost  cost;
      edge;
      mode mode;
      no-root-port;
      priority priority;
    }
    max-age seconds;
  }
    traceoptions {
      file filename <files number > <size size> <no-stamp | world-readable |
        no-world-readable>;
      flag flag;
    }
  vstp {
    bpdu-block-on-edge;
    disable;
    force-version stp;
    vlan (all | vlan-id | vlan-name) {
      bridge-priority priority;
      forward-delay seconds;
      hello-time seconds;
      interface (all | interface-name) {
        bpdu-timeout-action {
          alarm;
          block;
        }
        cost  cost;
        disable;
        edge;
        mode mode;
        no-root-port;
        priority priority;
      }
      max-age seconds;
      traceoptions {
        file filename <files number > <size size> <no-stamp | world-readable |
          no-world-readable>;
        flag flag;
      }
    }
  }
```

}

Related Documentation

- 802.1X for J-EX Series Switches Overview on page 2253
- Example: Configure Automatic VLAN Administration Using GVRP on page 1087
- Understanding MAC RADIUS Authentication on J-EX Series Switches
- Understanding Server Fail Fallback and 802.1X Authentication on J-EX Series Switches on page 2258
- IGMP Snooping on J-EX Series Switches Overview on page 2047
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261
- Understanding MSTP for J-EX Series Switches on page 1277
- Understanding Multiple VLAN Registration Protocol (MVRP) on J-EX Series Switches on page 1054
- Understanding Ethernet OAM Connectivity Fault Management for a J-EX Series Switch on page 3463
- Understanding Ethernet OAM Link Fault Management for a J-EX Series Switch on page 3427
- Understanding RSTP for J-EX Series Switches on page 1276
- Understanding STP for J-EX Series Switches on page 1275
- Understanding How to Use sFlow Technology for Network Monitoring on a J-EX Series Switch on page 3283
- Understanding VSTP for J-EX Series Switches on page 1281

action

Syntax

```
action {
  syslog;
  link-down;
}
```

Hierarchy Level [edit protocols **oam ethernet link-fault-management**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Define the action or actions to be taken when the OAM link fault management (LFM) fault event occurs.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 3431

action-profile

Syntax

```
action-profile profile-name;
  action {
    syslog;
    link-down;
  }
  event {
    link-adjacency-loss;
    link-event-rate;
    frame-error count;
    frame-period count;
    frame-period-summary count;
    symbol-period count;
  }
}
```

Hierarchy Level [edit protocols **oam ethernet link-fault-management**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure an Ethernet OAM link fault management (LFM) action profile by specifying a profile name.

The remaining statements are explained separately.

Options *profile-name*—Name of the action profile.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 3431

allow-remote-loopback

Syntax allow-remote-loopback;

Hierarchy Level [edit protocols **oam ethernet link-fault-management interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Advertise that the interface is capable of getting into loopback mode. Enable remote loopback in Ethernet OAM link fault management (LFM) on all Ethernet interfaces or the specified interface on the J-EX Series switch.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Ethernet OAM Link Fault Management on J-EX Series Switches on page 3428
- Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 3431

ethernet

Syntax

```
ethernet {
  connectivity-fault-management {
    action-profile profile-name {
      default-actions {
        interface-down;
      }
    }
    linktrace {
      age (30m | 10m | 1m | 30s | 10s);
      path-database-size path-database-size;
    }
    maintenance-domain domain-name {
      level number;
      mip-half-function (none | default |explicit);
      name-format (character-string | none | dns | mac+2oct);
      maintenance-association ma-name {
        continuity-check {
          hold-interval minutes;
          interval (10m | 10s | 1m | 1s| 100ms);
          loss-threshold number;
        }
        mep mep-id {
          auto-discovery;
          direction down;
          interface interface-name;
          remote-mep mep-id {
            action-profile profile-name;
          }
        }
      }
    }
  }
  link-fault-management{
    action-profile profile-name;
      action {
        syslog;
        link-down;
      }
      event {
        link-adjacency-loss;
        link-event-rate;
        frame-error count;
        frame-period count;
        frame-period-summary count;
        symbol-period count;
      }
      interface interface-name {
        link-discovery (active | passive);
        pdu-interval interval;
        event-thresholds threshold-value;
        remote-loopback;
        event-thresholds {
```

```
            frame-error count;
            frame-period count;
            frame-period-summary count;
            symbol-period count;
          }
        }
        negotiation-options {
          allow-remote-loopback;
          no-allow-link-events;
        }
    }
}
```

Hierarchy Level [edit protocols oam]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches. **connectivity-fault-management** introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Provide IEEE 802.3ah Operation, Administration, and Maintenance (OAM) support for Ethernet interfaces on J-EX Series switches or configure connectivity fault management (CFM) for IEEE 802.1ag Operation, Administration, and Management (OAM) support on the switches.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Ethernet OAM Link Fault Management on J-EX Series Switches on page 3428
- Example: Configuring Ethernet OAM Connectivity Fault Management on J-EX Series Switches on page 3465
- Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 3431
- Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 3468

event

Syntax

```
event {
  link-adjacency-loss;
  link-event-rate {
  frame-error count;
  frame-period count;
  frame-period-summary count;
  symbol-period count;
  }
}
```

Hierarchy Level [edit protocols **oam ethernet link-fault-management action-profile** *profile-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure link events in an action profile for Ethernet OAM link fault management (LFM).

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 3431

event-thresholds

Syntax

```
event-thresholds {
  frame-error count;
  frame-period count;
  frame-period-summary count;
  symbol-period count;
}
```

Hierarchy Level [edit protocols **oam ethernet link-fault-management interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure threshold limit values for link events in periodic OAM PDUs.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 3431

frame-error

Syntax frame-error *count*;

Hierarchy Level [edit protocols **oam ethernet link-fault-management event link-event-rate**],
[edit protocols **oam ethernet link-fault-management interface** *interface-name* **event-thresholds**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the threshold value for sending frame error events or taking the action specified in the action profile.

Frame errors occur on the underlying physical layer. The threshold is reached when the number of frame errors reaches the configured value.

Options *count*—Threshold count in seconds for frame error events.

Range: 1 through 100 seconds

Default: 1 second

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 3431

frame-period

Syntax frame-period *count*;

Hierarchy Level [edit protocols **oam ethernet link-fault-management event link-event-rate**],
[edit protocols **oam ethernet link-fault-management interface** *interface-name* **event-thresholds**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the number of frame errors within the last N frames that has exceeded a threshold.

Frame errors occur on the underlying physical layer. The threshold is reached when the number of frame errors reaches the configured value.

Options *count*—Threshold count in seconds for frame error events.

Range: 1 through 100 seconds

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 3431

frame-period-summary

Syntax frame-period-summary *count*;

Hierarchy Level [edit protocols **oam ethernet link-fault-management event link-event-rate**],
[edit protocols **oam ethernet link-fault-management interface** *interface-name* **event-thresholds**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the threshold value for sending frame period summary error events or taking the action specified in the action profile.

An errored frame second is any 1-second period that has at least one errored frame. This event is generated if the number of errored frame seconds is equal to or greater than the specified threshold for that period.

Options *count*—Threshold count in seconds for frame period summary error events.

Range: 1 through 100 seconds

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 3431

interface

Syntax

```
interface interface-name {
  link-discovery (active | passive);
  pdu-interval interval;
  event-thresholds threshold-value;
  remote-loopback;
  event-thresholds {
    frame-error count;
    frame-period count;
    frame-period-summary count;
    symbol-period count;
  }
  negotiation-options {
    allow-remote-loopback;
    no-allow-link-events;
  }
}
```

Hierarchy Level [edit protocols oam ethernet link-fault-management]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure Ethernet OAM link fault management (LFM) for all interfaces or for specific interfaces.

The remaining statements are explained separately.

Options *interface-name*—Name of the interface to be enabled for IEEE 802.3ah OAM link fault management (LFM) support.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Ethernet OAM Link Fault Management on J-EX Series Switches on page 3428
- Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 3431

link-adjacency-loss

Syntax link-adjacency-loss;

Hierarchy Level [edit protocols **oam ethernet link-fault-management action-profile event**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure **loss of adjacency** event with the IEEE 802.3ah link fault management (LFM) peer. When included, the loss of adjacency event triggers the action specified under the **action** statement.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Ethernet OAM Link Fault Management on J-EX Series Switches on page 3428
- Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 3431

link-discovery

Syntax link-discovery (active | passive);

Hierarchy Level [edit protocols **oam ethernet link-fault-management interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the discovery mode used for IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) support. The discovery process is triggered automatically when OAM 802.3ah functionality is enabled on an interface. Link monitoring is done when the interface sends periodic OAM PDUs.

Options *active*—In active mode, the interface discovers and monitors the peer on the link if the peer also supports IEEE 802.3ah OAM functionality.

passive—In passive mode, the peer initiates the discovery process.

Once the discovery process is initiated, both sides participate in discovery.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 3431

link-down

Syntax link-down;

Hierarchy Level [edit protocols **oam ethernet link-fault-management action-profile action**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Mark the interface as down for transit traffic.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 3431

link-event-rate

Syntax
```
link-event-rate {
  frame-error count;
  frame-period count;
  frame-period-summary count;
  symbol-period count;
}
```

Hierarchy Level [edit protocols **oam ethernet link-fault-management action-profile event**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the number of link fault management (LFM) events per second.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 3431

link-fault-management

Syntax

```
link-fault-management {
  action-profile profile-name;
    action {
      syslog;
      link-down;
    }
    event {
      link-adjacency-loss;
      link-event-rate;
      frame-error count;
      frame-period count;
      frame-period-summary count;
      symbol-period count;
    }
    interface interface-name {
      link-discovery (active | passive);
      pdu-interval interval;
      event-thresholds threshold-value;
      remote-loopback;
      event-thresholds {
        frame-error count;
        frame-period count;
        frame-period-summary count;
        symbol-period count;
      }
    }
    negotiation-options {
      allow-remote-loopback;
      no-allow-link-events;
    }
}
```

Hierarchy Level [edit protocols oam ethernet]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure Ethernet OAM link fault management (LFM) for all interfaces or for specific interfaces.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Ethernet OAM Link Fault Management on J-EX Series Switches on page 3428
- Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 3431

negotiation-options

Syntax

```
negotiation-options {
    allow-remote-loopback;
    no-allow-link-events;
}
```

Hierarchy Level [edit protocols **oam ethernet link-fault-management interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Enable and disable IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) features for Ethernet interfaces.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 3431

no-allow-link-events

Syntax no-allow-link-events;

Hierarchy Level [edit protocols **oam ethernet link-fault-management interface** *interface-name* **negotiation-options**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Disable the sending of link event TLVs.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 3431

oam

```
Syntax    oam {
            ethernet {
              connectivity-fault-management {
                action-profile profile-name {
                  default-actions {
                    interface-down;
                  }
                }
                linktrace {
                  age (30m | 10m | 1m | 30s | 10s);
                  path-database-size path-database-size;
                }
                maintenance-domain domain-name {
                  level number;
                  mip-half-function (none | default |explicit);
                  name-format (character-string | none | dns | mac+2oct);
                  maintenance-association ma-name {
                    continuity-check {
                      hold-interval minutes;
                      interval (10m | 10s | 1m | 1s| 100ms);
                      loss-threshold number;
                    }
                    mep mep-id {
                      auto-discovery;
                      direction down;
                      interface interface-name;
                      remote-mep mep-id {
                        action-profile profile-name;
                      }
                    }
                  }
                }
              }
              link-fault-management {
                action-profile profile-name;
                  action {
                    syslog;
                    link-down;
                  }
                  event {
                    link-adjacency-loss;
                    link-event-rate;
                    frame-error count;
                    frame-period count;
                    frame-period-summary count;
                    symbol-period count;
                  }
                  interface interface-name {
                    link-discovery (active | passive);
                    pdu-interval interval;
                    event-thresholds threshold-value;
                    remote-loopback;
```

```
            event-thresholds {
              frame-error count;
              frame-period count;
              frame-period-summary count;
              symbol-period count;
            }
          }
          negotiation-options {
            allow-remote-loopback;
            no-allow-link-events;
          }
        }
      }
```

Hierarchy Level [edit protocols]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches. **connectivity-fault-management** introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Provide IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) support for Ethernet interfaces on J-EX Series switches or configure connectivity fault management (CFM) for IEEE 802.1ag Operation, Administration, and Management (OAM) support on the switches.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Ethernet OAM Link Fault Management on J-EX Series Switches on page 3428
- Example: Configuring Ethernet OAM Connectivity Fault Management on J-EX Series Switches on page 3465
- Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 3431
- Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 3468

pdu-interval

Syntax pdu-interval *interval*;

Hierarchy Level [edit protocols **oam ethernet link-fault-management interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the periodic OAM PDU sending interval for fault detection. It is used for IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) support.

Options *interval*—Periodic OAM PDU sending interval.

Range: 400 through 1000 milliseconds

Default: 1000 milliseconds

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Ethernet OAM Link Fault Management on J-EX Series Switches on page 3428
- Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 3431

pdu-threshold

Syntax pdu-threshold *threshold-value*;

Hierarchy Level [edit protocols **oam ethernet link-fault-management interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure how many protocol data units (PDUs) are missed before declaring the peer lost in Ethernet OAM link fault management (LFM) for all interfaces or for specific interfaces.

Options ***threshold-value*** —Number of PDUs missed before declaring the peer lost.

Range: 3 through 10 PDUs

Default: 3 PDUs

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 3431

remote-loopback

Syntax remote-loopback;

Hierarchy Level [edit protocols **oam ethernet link-fault-management interface** *interface-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Set the data terminal equipment (DTE) in loopback mode. Remove the statement from the configuration to take the DTE out of loopback mode. It is used for IEEE 802.3ah Operation, Administration, and Maintenance (OAM) link fault management (LFM) support.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring Ethernet OAM Link Fault Management on J-EX Series Switches on page 3428
- Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 3431

symbol-period

Syntax symbol-period *count*;

Hierarchy Level [edit protocols **oam ethernet link-fault-management action-profile** *profile-name*; **event link-event-rate**] ,
[edit protocols **oam ethernet link-fault-management interface** *interface-name* **event-thresholds**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure the threshold for sending symbol period events or taking the action specified in the action profile.

Symbol code errors occur on the underlying physical layer. The symbol period threshold is reached when the number of symbol errors reaches the configured value within the period. You cannot configure the default value to a different value.

Options *count*—Threshold count in seconds for symbol period events.

Range: 1 through 100 seconds

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 3431

syslog

Syntax syslog;

Hierarchy Level [edit protocols **oam ethernet link-fault-management action-profile** *profile-name*; **action**]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Generate a system log message for the Ethernet Operation, Administration, and Maintenance (OAM) link fault management (LFM) event.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 3431

Operational Mode Commands for Ethernet OAM Link Fault Management

show oam ethernet link-fault-management

Syntax
show oam ethernet link-fault-management
<brief | detail>
<interface-name>

Release Information
Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description
Displays Operation, Administration, and Maintenance (OAM) link fault management (LFM) information for Ethernet interfaces.

Options
brief | detail—(Optional) Display the specified level of output.

interface-name —(Optional) Display link fault management information for the specified Ethernet interface only.

Required Privilege Level
view

Related Documentation
- Example: Configuring Ethernet OAM Link Fault Management on J-EX Series Switches on page 3428
- Configuring Ethernet OAM Link Fault Management (CLI Procedure) on page 3431

List of Sample Output
show oam ethernet link-fault-management brief on page 3462
show oam ethernet link-fault-management detail on page 3462

Output Fields
Table 455 on page 3458 lists the output fields for the **show oam ethernet link-fault-management** command. Output fields are listed in the approximate order in which they appear.

Table 455: show oam ethernet link-fault-management Output Fields

Field Name	Field Description	Level of Output
Status	Indicates the status of the established link. • **Fail**—A link fault condition exists. • **Running**—A link fault condition does not exist.	All levels
Discovery state	State of the discovery mechanism: • **Passive Wait** • **Send Any** • **Send Local Remote** • **Send Local Remote Ok**	All levels
Peer address	Address of the OAM peer.	All levels

Table 455: show oam ethernet link-fault-management Output Fields *(continued)*

Field Name	Field Description	Level of Output
Flags	Information about the interface. • **Remote-Stable**—Indicates remote OAM client acknowledgment of, and satisfaction with local OAM state information. **False** indicates that remote DTE has either not seen or is unsatisfied with local state information. **True** indicates that remote DTE has seen and is satisfied with local state information. • **Local-Stable**—Indicates local OAM client acknowledgment of, and satisfaction with remote OAM state information. **False** indicates that local DTE either has not seen or is unsatisfied with remote state information. **True** indicates that local DTE has seen and is satisfied with remote state information. • **Remote-State-Valid**—Indicates the OAM client has received remote state information found within Local Information TLVs of received Information OAM PDUs. **False** indicates that OAM client has not seen remote state information. **True** indicates that the OAM client has seen remote state information.	All levels
Remote loopback status	Indicates the remote loopback status. An OAM entity can put its remote peer into loopback mode using the Loopback control OAM PDU. In loopback mode, every frame received is transmitted back on the same port (except for OAM PDUs, which are needed to maintain the OAM session).	All levels
Remote entity information	Remote entity information. • **Remote MUX action**—Indicates the state of the multiplexer functions of the OAM sublayer. Device is forwarding non-OAM PDUs to the lower sublayer or discarding non-OAM PDUs. • **Remote parser action**—Indicates the state of the parser function of the OAM sublayer. Device is forwarding non-OAM PDUs to higher sublayer, looping back non-OAM PDUs to the lower sublayer, or discarding non-OAM PDUs. • **Discovery mode**—Indicates whether discovery mode is active or inactive. • **Unidirectional mode**—Indicates the ability to operate a link in a unidirectional mode for diagnostic purposes. • **Remote loopback mode**—Indicates whether remote loopback is supported or not supported. • **Link events**—Indicates whether interpreting link events is supported or not supported on the remote peer. • **Variable requests**—Indicates whether variable requests are supported or not supported. The Variable Request OAM PDU, is used to request one or more MIB variables from the remote peer.	All levels
OAM Receive Statistics		
Information	The number of information PDUs received.	**detail**
Event	The number of loopback control PDUs received.	**detail**
Variable request	The number of variable request PDUs received.	**detail**
Variable response	The number of variable response PDUs received.	**detail**
Loopback control	The number of loopback control PDUs received.	**detail**

Table 455: show oam ethernet link-fault-management Output Fields *(continued)*

Field Name	Field Description	Level of Output
Organization specific	The number of vendor organization specific PDUs received.	**detail**
OAM Transmit Statistics		
Information	The number of information PDUs transmitted.	**detail**
Event	The number of event notification PDUs transmitted.	**detail**
Variable request	The number of variable request PDUs transmitted.	**detail**
Variable response	The number of variable response PDUs transmitted.	**detail**
Loopback control	The number of loopback control PDUs transmitted.	**detail**
Organization specific	The number of vendor organization specific PDUs transmitted.	**detail**
OAM Received Symbol Error Event information		
Events	The number of symbol error event TLVs that have been received after the OAM sublayer was reset.	**detail**
Window	The symbol error event window in the received PDU. The protocol default value is the number of symbols that can be received in one second on the underlying physical layer.	**detail**
Threshold	The number of errored symbols in the period required for the event to be generated.	**detail**
Errors in period	The number of symbol errors in the period reported in the received event PDU.	**detail**
Total errors	The number of errored symbols that have been reported in received event TLVs after the OAM sublayer was reset. Symbol errors are coding symbol errors.	**detail**
OAM Received Frame Error Event Information		
Events	The number of errored frame event TLVs that have been received after the OAM sublayer was reset.	**detail**
Window	The duration of the window in terms of the number of 100 ms period intervals.	**detail**
Threshold	The number of detected errored frames required for the event to be generated.	**detail**
Errors in period	The number of detected errored frames in the period.	**detail**

Table 455: show oam ethernet link-fault-management Output Fields *(continued)*

Field Name	Field Description	Level of Output
Total errors	The number of errored frames that have been reported in received event TLVs after the OAM sublayer was reset. A frame error is any frame error on the underlying physical layer.	**detail**
OAM Received Frame Period Error Event Information		
Events	The number of frame seconds errors event TLVs that have been received after the OAM sublayer was reset.	**detail**
Window	The duration of the frame seconds window.	**detail**
Threshold	The number of frame seconds errors in the period.	**detail**
Errors in period	The number of frame seconds errors in the period.	**detail**
Total errors	The number of frame seconds errors that have been reported in received event TLVs after the OAM sublayer was reset.	**detail**
OAM Transmitted Symbol Error Event Information		
Events	The number of symbol error event TLVs that have been transmitted after the OAM sublayer was reset.	**detail**
Window	The symbol error event window in the transmitted PDU.	**detail**
Threshold	The number of errored symbols in the period required for the event to be generated.	**detail**
Errors in period	The number of symbol errors in the period reported in the transmitted event PDU.	**detail**
Total errors	The number of errored symbols reported in event TLVs that have been transmitted after the OAM sublayer was reset.	**detail**
OAM Transmitted Frame Error Event Information		
Events	The number of errored frame event TLVs that have been transmitted after the OAM sublayer was reset.	**detail**
Window	The duration of the window in terms of the number of 100 ms period intervals.	**detail**
Threshold	The number of detected errored frames required for the event to be generated.	**detail**
Errors in period	The number of detected errored frames in the period.	**detail**
Total errors	The number of errored frames that have been detected after the OAM sublayer was reset.	**detail**

show oam ethernet link-fault-management brief

```
user@host> show oam ethernet link-fault-management brief
Interface: ge-0/0/1
  Status: Running, Discovery state: Send Any
  Peer address: 00:90:69:72:2c:83
  Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
  Remote loopback status: Disabled on local port, Enabled on peer port
  Remote entity information:
    Remote MUX action: discarding, Remote parser action: loopback
    Discovery mode: active, Unidirectional mode: unsupported
    Remote loopback mode: supported, Link events: supported
    Variable requests: unsupported
```

show oam ethernet link-fault-management detail

```
user@host> show oam ethernet link-fault-management detail
  Interface: ge-0/0/1
    Status: Running, Discovery state: Send Any
    Peer address: 00:90:69:0a:07:14
    Flags:Remote-Stable Remote-State-Valid Local-Stable 0x50
    OAM receive statistics:
      Information: 186365, Event: 0, Variable request: 0, Variable response: 0
      Loopback control: 0, Organization specific: 0
    OAM transmit statistics:
      Information: 186347, Event: 0, Variable request: 0, Variable response: 0
      Loopback control: 0, Organization specific: 0
    OAM received symbol error event information:
      Events: 0, Window: 0, Threshold: 0
      Errors in period: 0, Total errors: 0
    OAM received frame error event information:
      Events: 0, Window: 0, Threshold: 0
      Errors in period: 0, Total errors: 0
    OAM received frame period error event information:
      Events: 0, Window: 0, Threshold: 0
      Errors in period: 0, Total errors: 0
    OAM transmitted symbol error event information:
      Events: 0, Window: 0, Threshold: 1
      Errors in period: 0, Total errors: 0
    OAM transmitted frame error event information:
      Events: 0, Window: 0, Threshold: 1
      Errors in period: 0, Total errors: 0
    Remote entity information:
      Remote MUX action: forwarding, Remote parser action: forwarding
      Discovery mode: active, Unidirectional mode: unsupported
      Remote loopback mode: supported, Link events: supported
      Variable requests: unsupported
```

CHAPTER 131

Ethernet OAM Connectivity Fault Management

Ethernet OAM Connectivity Fault Management—Overview

Understanding Ethernet OAM Connectivity Fault Management for a J-EX Series Switch

Ethernet interfaces on J-EX Series Switches and the Junos OS for J-EX Series switches support the IEEE 802.1ag standard for Operation, Administration, and Management (OAM). The IEEE 802.1ag specification provides for Ethernet connectivity fault management (CFM). CFM monitors Ethernet networks that might comprise one or more service instances for network-compromising connectivity faults.

The major features of CFM are:

- Fault monitoring using the continuity check protocol. This is a neighbor discovery and health check protocol that discovers and maintains adjacencies at the VLAN or link level.
- Path discovery and fault verification using the linktrace protocol.
- Fault isolation using the loopback protocol.

CFM partitions the service network into various administrative domains. For example, operators, providers, and customers might be part of different administrative domains. Each administrative domain is mapped into one maintenance domain providing enough information to perform its own management, thus avoiding security breaches and making end-to-end monitoring possible.

In a CFM maintenance domain, each service instance is called a maintenance association. A maintenance association can be thought of as a full mesh of maintenance association endpoints (MEPs) having similar characteristics. MEPs are active CFM entities generating and responding to CFM protocol messages. There is also a maintenance intermediate point (MIP), which is a CFM entity similar to the MEP, but more passive (MIPs only respond to CFM messages).

Each maintenance domain is associated with a maintenance domain level from 0 through 7. Level allocation is based on the network hierarchy, where outer domains are assigned a higher level than the inner domains. Configure customer end points to have the highest maintenance domain level. The maintenance domain level is a mandatory parameter that indicates the nesting relationships between various maintenance domains. The level is embedded in each CFM frame. CFM messages within a given level are processed by MEPs at that same level.

To enable CFM on an Ethernet interface, you must configure maintenance domains, maintenance associations, and maintenance association end points (MEPs). Figure 88 on page 3464 shows the relationships among maintenance domains, maintenance association end points (MEPs), and maintenance intermediate points (MIPs) configured on a switch.

Figure 88: Relationship Among MEPs, MIPs, and Maintenance Domain Levels

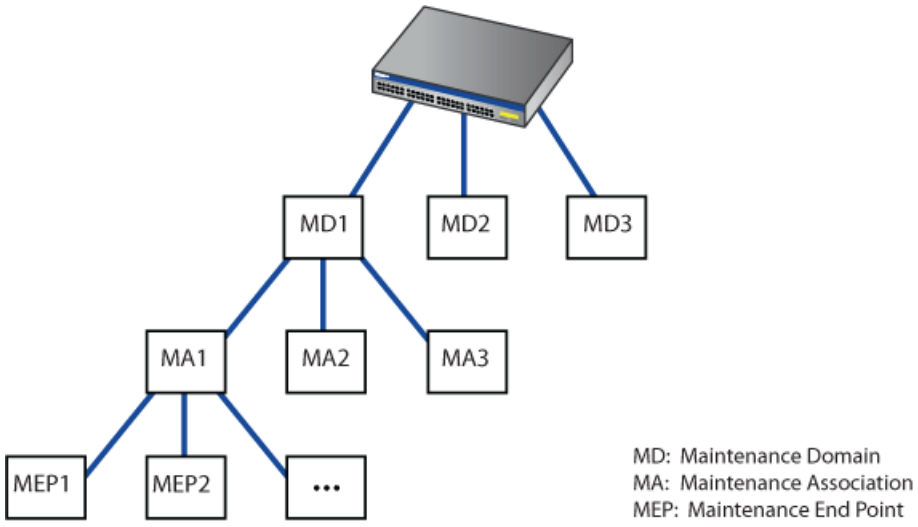

Related Documentation

- Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 3468
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

Example of Ethernet OAM Connectivity Fault Management Configuration

- Example: Configuring Ethernet OAM Connectivity Fault Management on J-EX Series Switches on page 3465

Example: Configuring Ethernet OAM Connectivity Fault Management on J-EX Series Switches

Ethernet interfaces on J-EX Series Switches and Junos OS for J-EX Series switches support the IEEE 802.1ag standard for Operation, Administration, and Management (OAM). The IEEE 802.1ag specification provides for Ethernet connectivity fault management (CFM).

This example describes how to enable and configure OAM CFM on a Gigabit Ethernet interface:

Requirements

This example uses the following hardware and software components:

- Two J-EX Series switches connected by a point-to-point Gigabit Ethernet link

Overview and Topology

CFM can be used to monitor the physical link between two switches. In the following example, two switches are connected by a point-to-point Gigabit Ethernet link. The link between these two switches is monitored using CFM.

Configuring Ethernet OAM Connectivity Fault Management on Switch 1

CLI Quick Configuration

To quickly configure Ethernet OAM CFM, copy the following commands and paste them into the switch terminal window:

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain]
set name-format character-string
set maintenance-domain private level 0
set maintenance-association private-ma
set continuity-check hold-interval 1s
```

Step-by-Step Procedure

To enable and configure OAM CFM on switch 1:

1. Specify the maintenance domain name format:

   ```
   [edit protocols oam ethernet connectivity-fault-management
   maintenance-domain]
   user@switch1# set name-format character-string
   ```

2. Specify the maintenance domain name and the maintenance domain level:

   ```
   [edit protocols oam ethernet connectivity-fault-management]
   user@switch1# set maintenance-domain private level 0
   ```

3. Create a maintenance association:

   ```
   [edit protocols oam ethernet connectivity-fault-management
   maintenance-domain private]
   user@switch1# set maintenance-association private-ma
   ```

4. Enable the continuity check protocol and specify the continuity check hold interval:

```
[edit protocols oam ethernet connectivity-fault-management
maintenance-domain private maintenance-association private-ma]
user@switch1# set continuity-check hold-interval 1s
```

5. Configure the maintenance association end point (MEP):

```
[edit protocols oam ethernet connectivity-fault-management
maintenance-domain private maintenance-association private-ma]
user@switch1# set mep 100 interface ge-1/0/1 auto-discovery direction down
```

Results Check the results of the configuration.

```
[edit]
user@switch1# show

  protocols {
    oam {
      ethernet {
        connectivity-fault-management {
          maintenance-domain private {
            level 0;
            maintenance-association private-ma {
              continuity-check {
                interval 1s;
              }
              mep 100 {
                interface ge-1/0/1;
                auto-discovery;
                direction down;
              }
            }
          }
        }
      }
    }
```

Configuring Ethernet OAM Connectivity Fault Management on Switch 2

CLI Quick Configuration To quickly configure Ethernet OAM CFM, copy the following commands and paste them into the switch terminal window:

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain]
set name-format character-string
set maintenance-domain private level 0
set maintenance-association private-ma
set continuity-check hold-interval 1s
```

Step-by-Step Procedure The configuration on switch 2 mirrors that on switch 2.

1. Specify the maintenance domain name format:

```
[edit protocols oam ethernet connectivity-fault-management]
user@switch2# set name-format character-string
```

2. Specify the maintenance domain name and the maintenance domain level:

```
[edit protocols oam ethernet connectivity-fault-management]
user@switch2# set maintenance-domain private level 0
```

3. Create a maintenance association:

```
[edit protocols oam ethernet connectivity-fault-management
maintenance-domain private]
user@switch2# set maintenance-association private-ma
```

4. Enable the continuity check protocol and specify the continuity check hold interval:

```
[edit protocols oam ethernet connectivity-fault-management
maintenance-domain private maintenance-association private-ma]
user@switch2# set continuity-check hold-interval 1s
```

5. Configure the maintenance association end point (MEP)

```
[edit protocols oam ethernet connectivity-fault-management
maintenance-domain private maintenance-association private-ma]
user@switch2# set mep 100 interface ge-0/2/5 auto-discovery direction down
```

Results Check the results of the configuration.

```
[edit]
user@switch2# show
```

```
protocols {
    oam {
        ethernet {
            connectivity-fault-management {
                maintenance-domain private {
                    level 0;
                    maintenance-association private-ma {
                        continuity-check {
                            interval 1s;
                        }
                        mep 100 {
                            interface ge-0/2/5;
                            auto-discovery;
                            direction down;
                        }
                    }
                }
            }
        }
    }
```

Verification

To confirm that the configuration is working properly, perform these tasks:

- Verifying That OAM CFM Has Been Configured Properly on page 3467

Verifying That OAM CFM Has Been Configured Properly

Purpose Verify that OAM CFM has been configured properly.

Action Use the show oam ethernet connectivity-fault-management interfaces detail command:

```
user@switch1# show oam ethernet connectivity-fault-management interfaces  detail
```

Sample Output

```
Interface name: ge-1/0/1.0, Interface status: Active, Link status: Up
  Maintenance domain name: private, Format: string, Level: 0
    Maintenance association name: private-ma, Format: string
```

```
Continuity-check status: enabled, Interval: 1ms, Loss-threshold: 3 frames
MEP identifier: 100, Direction: down, MAC address: 00:90:69:0b:4b:94
MEP status: running
Defects:
  Remote MEP not receiving CCM                 : no
  Erroneous CCM received                       : yes
  Cross-connect CCM received                   : no
  RDI sent by some MEP                         : yes
Statistics:
  CCMs sent                                    : 76
  CCMs received out of sequence                : 0
  LBMs sent                                    : 0
  Valid in-order LBRs received                 : 0
  Valid out-of-order LBRs received             : 0
  LBRs received with corrupted data            : 0
  LBRs sent                                    : 0
  LTMs sent                                    : 0
  LTMs received                                : 0
  LTRs sent                                    : 0
  LTRs received                                : 0
  Sequence number of next LTM request          : 0
Remote MEP count: 2
  Identifier    MAC address        State    Interface
   2001     00:90:69:0b:7f:71       ok    ge-0/2/5.0
```

Meaning When the output displays continuity-check status is **enabled** and displays details of the remote MEP, it means that connectivity fault management (CFM) has been configured properly.

Related Documentation

- Understanding Ethernet OAM Connectivity Fault Management for a J-EX Series Switch on page 3463
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

Configuring Ethernet OAM Connectivity Fault Management

- Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 3468

Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure)

Ethernet interfaces on J-EX Series Switches and the Junos OS for J-EX Series switches support the IEEE 802.1ag standard for Operation, Administration, and Management (OAM). The IEEE 802.1ag specification provides for Ethernet connectivity fault management (CFM).

This topic describes these tasks:

1. Creating the Maintenance Domain on page 3469
2. Configuring the Maintenance Domain MIP Half Function on page 3469
3. Creating a Maintenance Association on page 3470
4. Configuring the Continuity Check Protocol on page 3470

5. Configuring a Maintenance Association End Point on page 3470
6. Configuring a Connectivity Fault Management Action Profile on page 3471
7. Configuring the Linktrace Protocol on page 3472

Creating the Maintenance Domain

A maintenance domain comprises network entities such as operators, providers, and customers. To enable connectivity fault management (CFM) on an Ethernet interface, you must create a maintenance domains, maintenance associations, and MEPs.

To create a maintenance domain:

1. Specify a name for the maintenance domain:

```
[edit protocols oam ethernet connectivity-fault-management]
user@switch# set maintenance-domain domain-name
```

2. Specify a format for the maintenance domain name. If you specify **none**, no name is configured:

   - A plain ASCII character string
   - A domain name service (DNS) format
   - A media access control (MAC) address plus a two-octet identifier in the range 0 through 65,535
   - **none**

```
[edit protocols oam ethernet connectivity-fault-management
maintenance-domain domain-name]
user@switch# set name-format format
```

   For example, to specify the name format as MAC address plus a two-octet identifier:

```
[edit protocols oam ethernet connectivity-fault-management
maintenance-domain domain-name]
user@switch# set name-format mac+2oct
```

3. Configure the maintenance domain level, which is used to indicate the nesting relationship between this domain and other domains. Use a value from 0 through 7:

```
[edit protocols oam ethernet connectivity-fault-management
maintenance-domain domain-name]
user@switch# set level level
```

Configuring the Maintenance Domain MIP Half Function

MIP Half Function (MHF) divides the maintenance association intermediate point (MIP) functionality into two unidirectional segments, improves visibility with minimal configuration, and improves network coverage by increasing the number of points that can be monitored. MHF extends monitoring capability by responding to loop-back and link-trace messages to help isolate faults. Whenever a MIP is configured, the MIP half function value for all maintenance domains and maintenance associations must be the same.

To configure the MIP half function:

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
domain-name]
user@switch# set mip-half-function (none | default | explicit)
```

Creating a Maintenance Association

In a CFM maintenance domain, each service instance is called a maintenance association.

To create a maintenance association:

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
domain-name]
user@switch# set maintenance-association ma-name
```

Configuring the Continuity Check Protocol

The continuity check protocol is used for fault detection by a maintenance association end point (MEP) within a maintenance association. The MEP periodically sends continuity check multicast messages. The receiving MEPs use the continuity check messages (CCMs) to build a MEP database of all MEPs in the maintenance association.

To configure the continuity check protocol:

1. Enable the continuity check protocol:

   ```
   [edit protocols oam ethernet connectivity-fault-management
   maintenance-domain domain-name maintenance-association ma-name]
   user@switch# set continuity-check
   ```

2. Specify the continuity check hold interval. The hold interval is the number of minutes to wait before flushing the MEP database if no updates occur. The default value is 10 minutes.

   ```
   [edit protocols oam ethernet connectivity-fault-management
   maintenance-domain domain-name maintenance-association ma-name
   continuity-check]
   user@switch# set hold-interval number
   ```

3. Specify the CCM interval. The interval is the time between the transmission of CCMs. You can specify 10 minutes (10m), 1 minute (1m), 10 seconds (10s), 1 second (1s), 100 milliseconds (100ms), or 10 milliseconds (10ms).

   ```
   [edit protocols oam ethernet connectivity-fault-management
   maintenance-domain domain-name maintenance-association ma-name
   continuity-check]
   user@switch# set interval number
   ```

4. Specify the number of CCMs (that is, protocol data units) that can be lost before the MEP is marked as down. The default number of protocol data units (PDUs) is 3.

   ```
   [edit protocols oam ethernet connectivity-fault-management
   maintenance-domain domain-name maintenance-association ma-name
   continuity-check]
   user@switch# set loss-threshold number
   ```

Configuring a Maintenance Association End Point

To configure a maintenance association end point:

1. Specify an ID for the MEP. The value can be from 1 through 8191.

```
[edit protocols oam ethernet connectivity-fault-management
maintenance-domain domain-name maintenance-association ma-name]
user@switch# set mep mep-id]
```

2. Enable maintenance endpoint automatic discovery if you want to have the MEP accept continuity check messages (CCMs) from all remote MEPs of the same maintenance association:

```
[edit protocols oam ethernet connectivity-fault-management
maintenance-domain domain-name maintenance-association ma-name mep mep-id
user@switch# set auto-discovery
```

3. You can specify that CFM packets (CCMs) be transmitted only in one direction for the MEP, that is, the direction be set as **down** so that CCMs are transmitted only out of (not into) the interface configured on this MEP.

```
[edit protocols oam ethernet connectivity-fault-management
maintenance-domain domain-name maintenance-association ma-name mep mep-id]
user@switch# set direction down
```

4. Specify the logical interface to which the MEP is attached. It can be either an access interface or a trunk interface. If you specify a trunk interface, the VLAN associated with that interface must have a VLAN ID.

```
[edit protocols oam ethernet connectivity-fault-management
maintenance-domain domain-name maintenance-association ma-name mep mep-id]
user@switch# set interface interface-name
```

5. You can configure a remote MEP from which CCMs are expected. If autodiscovery is not enabled, the remote MEP must be configured under the **mep** statement. If the remote MEP is not configured under the **mep** statement, the CCMs from the remote MEP are treated as errors.

```
[edit protocols oam ethernet connectivity-fault-management
maintenance-domain domain-name maintenance-association ma-name mep mep-id]
user@switch# set remote-mep mep-id
```

Configuring a Connectivity Fault Management Action Profile

You can configure an action profile and specify the action to be taken when any of the configured events occur. Alternatively, you can configure an action profile and specify default actions when connectivity to a remote MEP fails.

To configure an action profile:

1. Specify a name for an action profile:

```
[edit protocols oam ethernet connectivity-fault-management]
user@switch# set action-profile profile-name
```

2. Configure the action of the action profile:

```
[edit protocols oam ethernet connectivity-fault-management action-profile
profile-name]
user@switch# set action interface-down
```

3. Configure one or more events under the action profile, the occurrence of which will trigger the corresponding action to be taken:

```
[edit protocols oam ethernet connectivity-fault-management action-profile
profile-name]
user@switch# set event event
```

See the *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

Configuring the Linktrace Protocol

The linktrace protocol is used for path discovery between a pair of maintenance points. Linktrace messages are triggered by an administrator using the **traceroute** command to verify the path between a pair of MEPs under the same maintenance association. Linktrace messages can also be used to verify the path between a MEP and a MIP under the same maintenance domain.

To configure the linktrace protocol:

1. Configure the linktrace path age timer. If no response to a linktrace request is received, the request and response entries are deleted after the age timer expires:

```
[edit protocols oam ethernet connectivity-fault-management]
user@switch# set linktrace age time
```

2. Configure the number of linktrace reply entries to be stored per linktrace request:

```
[edit protocols oam ethernet connectivity-fault-management]
user@switch# set linktrace path-database-size path-database-size
```

Related Documentation

- Example: Configuring Ethernet OAM Connectivity Fault Management on J-EX Series Switches on page 3465
- Understanding Ethernet OAM Connectivity Fault Management for a J-EX Series Switch on page 3463
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

Configuration Statements for Ethernet OAM Connectivity Fault Management

- [edit protocols] Configuration Statement Hierarchy on page 3472

[edit protocols] Configuration Statement Hierarchy

```
protocols {
  connections {
    remote-interface-switch connection-name {
      interface interface-name.unit-number;
      transmit-lsp label-switched-path;
      receive-lsp label-switched-path;
    }
  }
  dot1x {
```

```
    authenticator {
      authentication-profile-name profile-name;
      interface (all | [ interface-names ]) {
        disable;
        guest-vlan ( vlan-id | vlan-name);
        mac-radius <restrict>;
        maximum-requests number;
        no-reauthentication;
        quiet-period seconds;
        reauthentication {
          interval seconds;
        }
        retries number;
        server-fail (deny | permit | use-cache | vlan-id | vlan-name);
        server-reject-vlan (vlan-id | vlan-name);
        server-timeout seconds;
        supplicant (multiple | single | single-secure);
        supplicant-timeout seconds;
        transmit-period seconds;
      }
      static mac-address {
        interface interface-name;
        vlan-assignment (vlan-id |vlan-name);
      }
  }
  gvrp {
    <enable | disable>;
    interface (all | [interface-name]) {
      disable;
    }
    join-timer millseconds;
    leave-timer milliseconds;
    leaveall-timer milliseconds;
  }
  igmp-snooping {
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>
        <match regex>;
      flag flag (detail | disable | receive | send);
    }
    vlan (vlan-id | vlan-number) {
      data-forwarding {
        source {
          groups group-prefix;
        }
        receiver {
          source-vlans vlan-list;
          install ;
        }
      }
      disable {
        interface interface-name
      }
      immediate-leave;
      interface interface-name {
        group-limit limit;
```

```
        multicast-router-interface;
        static {
          group ip-address;
        }
      }
      proxy ;
      query-interval seconds;
      query-last-member-interval seconds;
      query-response-interval seconds;
      robust-count number;
    }
  }
  lldp {
    disable;
    advertisement-interval seconds;
    hold-multiplier number;
    interface (all | interface-name) {
      disable;
    }
    traceoptions {
      file filename <files number> <size size> <world-readable | no-world-readable>
        <match regex>;
      flag flag (detail | disable | receive | send);
    }
  }
  lldp-med {
    disable;
    fast-start number;
    interface (all | interface-name) {
      disable;
      location {
        elin number;
        civic-based {
          what number;
          country-code code;
          ca-type {
            number {
              ca-value value;
            }
          }
        }
      }
    }
  }
  mpls {
    interface ( all | interface-name );
    label-switched-path lsp-name to remote-provider-edge-switch;
    path destination {
      <address | hostname> <strict | loose>
    }
  mstp {
    disable;
    bpdu-block-on-edge;
    bridge-priority priority;
    configuration-name name;
    forward-delay seconds;
```

```
    hello-time seconds;
    interface (all | interface-name) {
      disable;
      bpdu-timeout-action {
        block;
        alarm;
      }
      cost  cost;
      edge;
      mode mode;
      no-root-port;
      priority priority;
    }
    max-age seconds;
    max-hops hops;
    msti msti-id {
      vlan (vlan-id | vlan-name);
      interface interface-name {
        disable;
        cost  cost;
        edge;
        mode mode;
        priority priority;
      }
    }
    revision-level revision-level;
    traceoptions {
      file filename <files number > <size size> <no-stamp | world-readable |
        no-world-readable>;
      flag flag;
    }
  }
  mvrp {
    disable
    interface (all | interface-name) {
      disable;
      join-timer milliseconds;
      leave-timer milliseconds;
      leaveall-timer milliseconds;
      registration (forbidden | normal);
    }
    no-dynamic-vlan;
    traceoptions {
      file filename <files number > <size size> <no-stamp | world-readable |
        no-world-readable>;
      flag flag;
    }
  }
  oam {
    ethernet{
      connectivity-fault-management {
        action-profile profile-name {
          default-actions {
            interface-down;
          }
        }
```

```
    linktrace {
      age (30m | 10m | 1m | 30s | 10s);
      path-database-size path-database-size;
    }
    maintenance-domain domain-name {
      level number;
      mip-half-function (none | default |explicit);
      name-format (character-string | none | dns | mac+2oct);
      maintenance-association ma-name {
        continuity-check {
          hold-interval minutes;
          interval (10m | 10s | 1m | 1s| 100ms);
          loss-threshold number;
        }
        mep mep-id {
          auto-discovery;
          direction down;
          interface interface-name;
          remote-mep mep-id {
            action-profile profile-name;
          }
        }
      }
    }
  }
  link-fault-management {
    action-profile profile-name;
      action {
        syslog;
        link-down;
      }
      event {
        link-adjacency-loss;
        link-event-rate;
        frame-error count;
        frame-period count;
        frame-period-summary count;
        symbol-period count;
      }
      interface interface-name {
        link-discovery (active | passive);
        pdu-interval interval;
        event-thresholds threshold-value;
        remote-loopback;
        event-thresholds {
          frame-errorcount;
          frame-period count;
          frame-period-summary count;
          symbol-period count;
        }
      }
      negotiation-options {
        allow-remote-loopback;
        no-allow-link-events;
      }
  }
```

```
    }
  }
  rstp {
    disable;
    bpdu-block-on-edge;
    bridge-priority priority;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
      disable;
      bpdu-timeout-action {
        block;
        alarm;
      }
      cost  cost;
      edge;
      mode mode;
      no-root-port;
      priority priority;
    }
    max-age seconds;
    }
    traceoptions {
      file filename <files number > <size size> <no-stamp | world-readable |
        no-world-readable>;
      flag flag;
    }
  }
  sflow {
    agent-id
    collector {
      ip-address;
      udp-port port-number;
    }
    disable;
    interfaces interface-name {
      disable;
      polling-interval seconds;
      sample-rate number;
    }
    polling-interval seconds;
    sample-rate number;
    source-ip
  }
  stp {
    disable;
    bridge-priority priority;
    forward-delay seconds;
    hello-time seconds;
    interface (all | interface-name) {
      disable;
      bpdu-timeout-action {
        block;
        alarm;
      }
      cost  cost;
```

```
        edge;
        mode mode;
        no-root-port;
        priority priority;
      }
      max-age seconds;
    }
      traceoptions {
        file filename <files number > <size size> <no-stamp | world-readable |
          no-world-readable>;
        flag flag;
      }
    vstp {
      bpdu-block-on-edge;
      disable;
      force-version stp;
      vlan (all | vlan-id | vlan-name) {
        bridge-priority priority;
        forward-delay seconds;
        hello-time seconds;
        interface (all | interface-name) {
          bpdu-timeout-action {
            alarm;
            block;
          }
          cost  cost;
          disable;
          edge;
          mode mode;
          no-root-port;
          priority priority;
        }
        max-age seconds;
        traceoptions {
          file filename <files number > <size size> <no-stamp | world-readable |
            no-world-readable>;
          flag flag;
        }
      }
    }
  }
```

Related Documentation

- 802.1X for J-EX Series Switches Overview on page 2253
- Example: Configure Automatic VLAN Administration Using GVRP on page 1087
- Understanding MAC RADIUS Authentication on J-EX Series Switches
- Understanding Server Fail Fallback and 802.1X Authentication on J-EX Series Switches on page 2258
- IGMP Snooping on J-EX Series Switches Overview on page 2047
- Understanding 802.1X and LLDP and LLDP-MED on J-EX Series Switches on page 2261
- Understanding MSTP for J-EX Series Switches on page 1277

- Understanding Multiple VLAN Registration Protocol (MVRP) on J-EX Series Switches on page 1054
- Understanding Ethernet OAM Connectivity Fault Management for a J-EX Series Switch on page 3463
- Understanding Ethernet OAM Link Fault Management for a J-EX Series Switch on page 3427
- Understanding RSTP for J-EX Series Switches on page 1276
- Understanding STP for J-EX Series Switches on page 1275
- Understanding How to Use sFlow Technology for Network Monitoring on a J-EX Series Switch on page 3283
- Understanding VSTP for J-EX Series Switches on page 1281

action-profile (Applying to OAM CFM, for J-EX Series Switch Only)

Syntax

```
action-profile profile-name {
  default-actions {
    interface-down;
  }
}
```

Hierarchy Level [edit protocols **oam ethernet connectivity-fault-management**]

Release Information Statement introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Configure a name and default action for an action profile.

Options

profile-name—Name of the action profile.

default-actions—Defines the action to be taken when connectivity to the remote MEP is lost.

interface-down—Brings the interface down when a remote MEP connectivity failure is detected.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 3468
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

age (J-EX Series Switch Only)

Syntax age (30m | 10m | 1m | 30s | 10s);

Hierarchy Level [edit protocols **oam ethernet connectivity-fault-management linktrace**]

Release Information Statement introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Configure the time to wait (in minutes or seconds) for a response. If no response is received, the request and response entry is deleted from the linktrace database.

Default 10 minutes

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 3468
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

auto-discovery (J-EX Series Switch Only)

Syntax auto-discovery;

Hierarchy Level [edit protocols **oam ethernet connectivity-fault-management maintenance-domain** *domain-name* **maintenance-association** *ma-name* **mep** *mep-id*]

Release Information Statement introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Enable the MEP to accept continuity check messages from all remote MEPs.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 3468
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

connectivity-fault-management (J-EX Series Switch Only)

Syntax

```
connectivity-fault-management {
  action-profile profile-name {
    default-actions {
      interface-down;
    }
  }
  linktrace {
    age (30m | 10m | 1m | 30s | 10s);
    path-database-size path-database-size;
  }
  maintenance-domain domain-name {
    level number;
    mip-half-function (none | default |explicit);
    name-format (character-string | none | dns | mac+2oct);
    maintenance-association ma-name {
      continuity-check {
        hold-interval minutes;
        interval (10m | 10s | 1m | 1s| 100ms);
        loss-threshold number;
      }
      mep mep-id {
        auto-discovery;
        direction down;
        interface interface-name;
        remote-mep mep-id {
          action-profile profile-name;
        }
      }
    }
  }
}
```

Hierarchy Level [edit protocols oam ethernet]

Release Information Statement introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Configure connectivity fault management for IEEE 802.1ag Operation, Administration, and Management (OAM) support.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 3468
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

continuity-check (J-EX Series Switch Only)

Syntax

```
continuity-check {
  hold-interval minutes;
  interval (10m | 10s | 1m | 1s| 100ms);
  loss-threshold number;
}
```

Hierarchy Level [edit protocols **oam ethernet connectivity-fault-management maintenance-domain** *domain-name* **maintenance-association** *ma-name*]

Release Information Statement introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Specify continuity check protocol options.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 3468
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

direction (J-EX Series Switch Only)

Syntax direction down;

Hierarchy Level [edit protocols **oam ethernet connectivity-fault-management maintenance-domain** *domain-name* **maintenance-association** *ma-name* **mep** *mep-id*]

Release Information Statement introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Specify that connectivity fault management (CFM) packets (CCMs) be transmitted only in one direction for the MEP, that is, the direction be set as **down** so that CCMs are transmitted only out of (not into) the interface configured on this MEP.

Options **down**—Down MEP CCMs are transmitted only out (not into) of the interface configured on this MEP.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 3468
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

hold-interval (OAM CFM, for J-EX Series Switch Only)

Syntax hold-interval *minutes*;

Hierarchy Level [edit protocols **oam ethernet connectivity-fault-management maintenance-domain** *domain-name* **maintenance-association** *ma-name* **continuity-check**]

Release Information Statement introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Configure the time to wait before flushing the maintenance association end point (MEP) database, if no updates occur.

Options ***minutes***—Time to wait, in minutes.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 3468
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

interface (OAM CFM, for J-EX Series Switch Only)

Syntax interface (*interface-name* | ((ge- | xe-) (*fpc/pic/port* | *fpc/pic/port.unit-number* | *fpc/pic/port.unit-number* vlan *vlan-id*)));

Hierarchy Level [edit protocols **oam ethernet connectivity-fault-management maintenance-domain** *domain-name* **maintenance-association** *ma-name* **mep** *mep-id*]

Release Information Statement introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Configure IEEE 802.1ag Operation, Administration, and Management (OAM) Connectivity Fault Management (CFM) support for the specified interface.

Options ***interface-name***—Interface to which the MEP is attached. It can be a physical Ethernet interface or a logical interface. If the interface is a trunk interface, the VLAN associated with the interface must have a VLAN ID.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 3468
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

interval (J-EX Series Switch Only)

Syntax interval (10m | 10s | 1m | 1s | 100ms | 10ms);

Hierarchy Level [edit protocols **oam ethernet connectivity-fault-management maintenance-domain** *domain-name* **maintenance-association** *ma-name* **continuity-check**]

Release Information Statement introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Configure the time between continuity check messages.

Options

10m—10 minutes.

10s—10 seconds.

1m—1 minute.

1s—1 second.

100ms—100 milliseconds.

10ms—10 milliseconds.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 3468
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

level (J-EX Series Switch Only)

Syntax level *number*;

Hierarchy Level [edit protocols **oam ethernet connectivity-fault-management maintenance-domain** *domain-name*]

Release Information Statement introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Configure a number to be used in CFM messages to identify the maintenance association.

Options ***number***—Number used to identify the maintenance domain to which the CFM message belongs.

Range: 0 through 7

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 3468
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

linktrace (J-EX Series Switch Only)

Syntax
```
linktrace {
  age (30m | 10m | 1m | 30s | 10s);
  path-database-size path-database-size;
}
```

Hierarchy Level [edit protocols **oam ethernet connectivity-fault-management**]

Release Information Statement introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Configure connectivity fault management linktrace parameters.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 3468
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

loss-threshold (J-EX Series Switch Only)

Syntax loss-threshold *number*;

Hierarchy Level [edit protocols **oam ethernet connectivity-fault-management maintenance-domain** *domain-name* **maintenance-association** *ma-name* **continuity-check**]

Release Information Statement introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Configure the number of continuity check messages that can be lost before the remote MEP is marked as down.

Options ***number***—Number of continuity check messages that can be lost before the remote MEP is marked down.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 3468
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

maintenance-association (J-EX Series Switch Only)

Syntax

```
maintenance-association ma-name {
  continuity-check {
    hold-interval minutes;
    interval (10m | 10s | 1m | 1s| 100ms);
    loss-threshold number;
  }
  mep mep-id {
    auto-discovery;
    direction down;
    interface interface-name;
    remote-mep mep-id {
      action-profile profile-name;
    }
  }
}
```

Hierarchy Level [edit protocols **oam ethernet connectivity-fault-management maintenance-domain** *domain-name*]

Release Information Statement introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Configure the name of the maintenance association in IEEE-compliant format.

Options ***ma-name***—The name of the maintenance association within the maintenance domain.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 3468
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

maintenance-domain (J-EX Series Switch Only)

Syntax

```
maintenance-domain domain-name {
  level number;
  mip-half-function (none | default |explicit);
  name-format (character-string | none | dns | mac+2oct);
  maintenance-association ma-name {
    continuity-check {
      hold-interval minutes;
      interval (10m | 10s | 1m | 1s| 100ms);
      loss-threshold number;
    }
    mep mep-id {
      auto-discovery;
      direction down;
      interface interface-name;
      remote-mep mep-id {
        action-profile profile-name;
      }
    }
  }
}
```

Hierarchy Level [edit protocols **oam ethernet connectivity-fault-management**]

Release Information Statement introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Configure the name of the maintenance domain in IEEE-compliant format.

Options ***domain-name***—The name for the maintenance domain.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 3468
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

mep (J-EX Series Switch Only)

Syntax

```
mep mep-id {
  auto-discovery;
  direction down;
  interface interface-name;
  remote-mep mep-id {
    action-profile profile-name;
  }
}
```

Hierarchy Level [edit protocols **oam ethernet connectivity-fault-management maintenance-domain** *domain-name* **maintenance-association** *ma-name*]

Release Information Statement introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Configure the numeric identifier of the maintenance association end point (MEP) within the maintenance association.

Options ***mep-id***—Numeric identifier of the MEP.

Range: 1 through 8191

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 3468
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

mip-half-function (J-EX Series Switch Only)

Syntax mip-half-function (none | default | explicit);

Hierarchy Level [edit protocols **oam ethernet connectivity-fault-management maintenance-domain** *domain-name*]

Release Information Statement introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Specify the OAM Ethernet CFM maintenance domain MIP half functions.

NOTE: Whenever a MIP is configured, the MIP half function value for all maintenance domains and maintenance associations must be the same.

Options **none**—Specify to not use the mip-half-function.

default—Specify to use the default mip-half-function.

explicit—Specify an explicit mip-half-function.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 3468
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

name-format (J-EX Series Switch Only)

Syntax name-format (character-string | none | dns | mac+2oct);

Hierarchy Level [edit protocols **oam ethernet connectivity-fault-management maintenance-domain** *domain-name*]

Release Information Statement introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Specify the format of the maintenance domain name.

Options

character-string—The name is an ASCII character string.

none—Name format **none** means that maintenance domain name is not used.

dns—Name is in domain name service (DNS) format. For example: www.juniper.net.

mac+2oct—Name is the MAC address plus a two-octet maintenance association identifier. For example: 08:00:22:33:44:55.100.

Default: character-string

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 3468
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

path-database-size (J-EX Series Switch Only)

Syntax path-database-size *path-database-size*;

Hierarchy Level [edit protocols **oam ethernet connectivity-fault-management linktrace**]

Release Information Statement introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Specify the number of linktrace reply entries to be stored per linktrace request.

Options

path-database-size—Database size (number of entries stored per request).

Range: 1 through 500

Default: 100

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 3468
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

remote-mep (J-EX Series Switch Only)

Syntax

```
remote-mep mep-id {
    action-profile profile-name;
}
```

Hierarchy Level

[edit protocols **oam ethernet connectivity-fault-management maintenance-domain** *domain-name* **maintenance-association** *ma-name* **mep** *mep-id*]

Release Information

Statement introduced in Junos OS Release 10.2 for J-EX Series switches.

Description

Specify the numeric identifier of the remote maintenance association end point (MEP) within the maintenance association.

Options

mep-id—Specify the numeric identifier of the MEP.

Range: 1 through 8191

The remaining statement is explained separately.

Required Privilege Level

interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Ethernet OAM Connectivity Fault Management (CLI Procedure) on page 3468
- *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/

Operational Mode Commands for Ethernet OAM Connectivity Fault Management

clear oam ethernet connectivity-fault-management statistics

Syntax **clear oam ethernet connectivity-fault-management statistics**
<interface *ethernet-interface-name*>
<level *md-level*>

Release Information Command introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Clear all statistics maintained by CFM.

Options **interface *ethernet-interface-name***—(Optional) Clear CFM statistics only for MEPs attached to the specified Ethernet physical interface.

level *level*—(Optional) Clear CFM statistics only for MEPs within CFM maintenance domains (MDs) of the specified level.

Required Privilege Level clear

Related Documentation
- **show oam ethernet connectivity-fault-management interfaces on page 3498**
- **show oam ethernet connectivity-fault-management linktrace path-database on page 3504**
- **show oam ethernet connectivity-fault-management mip on page 3512**

List of Sample Output **clear oam ethernet connectivity-fault-management statistics on page 3493**

Output Fields When you enter this command, you are provided feedback on the status of your request.

clear oam ethernet connectivity-fault-management statistics

```
user@host> clear oam ethernet connectivity-fault-management statistics
Cleared statistics of all CFM sessions
```

show oam ethernet connectivity-fault-management forwarding-state

Syntax **show oam ethernet connectivity-fault-management forwarding-state <brief | detail | extensive>**

Release Information Command introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management forwarding state information for Ethernet interfaces.

Options interface *interface-name*—Display forwarding state information for the specified Ethernet interface only.

brief | detail | extensive—(Optional) Display the specified level of output.

Required Privilege Level view

Related Documentation

- **clear oam ethernet connectivity-fault-management statistics on page 3493**
- **show oam ethernet connectivity-fault-management linktrace path-database on page 3504**
- **show oam ethernet connectivity-fault-management mip on page 3512**

List of Sample Output

show oam ethernet connectivity-fault-management forwarding-state on page 3495
show oam ethernet connectivity-fault-management forwarding-state interface on page 3495
show oam ethernet connectivity-fault-management forwarding-state interface detail on page 3496
show oam ethernet connectivity-fault-management forwarding-state interface interface-name on page 3496

Output Fields Table 456 on page 3494 lists the output fields for the **show oam ethernet connectivity-fault-management forwarding-state** command. Output fields are listed in the approximate order in which they appear.

Table 456: show oam ethernet connectivity-fault-management forwarding-state Output Fields

Field Name	Field Description	Level of Output
Interface name	Interface identifier.	All levels
Filter action	Filter action for messages at the level.	All levels
Nexthop type	Next-hop type.	All levels
Nexthop index	Next-hop index number.	**brief**
Level	Maintenance domain (MD) level.	**detail**

Table 456: show oam ethernet connectivity-fault-management forwarding-state Output Fields *(continued)*

Field Name	Field Description	Level of Output
Direction	MEP direction configured.	none
CEs	Number of customer edge (CE) interfaces.	All levels

show oam ethernet connectivity-fault-management forwarding-state

```
user@host> show oam ethernet connectivity–fault-management forwarding-state
CEs: 3

  Maintenance domain forwarding state:

  Level    Direction      Filter action     Nexthop      Nexthop
                                            type         index
  0                       Drop              none
  1                       Drop              none
  2                       Drop              none
  3                       Drop              none
  4                       Drop              none
  5                       Drop              none
  6                       Drop              none
  7                       Drop              none
```

show oam ethernet connectivity-fault-management forwarding-state interface

```
user@host> show oam ethernet connectivity-fault-management forwarding-state interface
Interface name: ge-3/0/0.0
  Maintenance domain forwarding state:

  Level    Direction      Filter action     Nexthop      Nexthop
                                            type         index
  0                       Drop              none
  1                       Drop              none
  2                       Drop              none
  3                       Drop              none
  4                       Drop              none
  5                       Drop              none
  6                       Drop              none
  7        down           Receive           none

Interface name: xe-0/0/0.0
Instance name: __+bd1__
  Maintenance domain forwarding state:

  Level    Direction      Filter action     Nexthop      Nexthop
                                            type         index
  0                       Drop              none
  1                       Drop              none
  2                       Drop              none
  3                       Drop              none
  4                       Drop              none
  5                       Drop              none
  6                       Drop              none
  7        down           Receive           none
```

show oam ethernet connectivity-fault-management forwarding-state interface detail

```
user@host> show oam ethernet connectivity-fault-management forwarding-state interface
detail
Interface name: ge-3/0/0.0

  Level: 0
  Filter action: Drop
  Nexthop type: none

  Level: 1
  Filter action: Drop
  Nexthop type: none

  Level: 2
  Filter action: Drop
  Nexthop type: none

  Level: 3
  Filter action: Drop
  Nexthop type: none

  Level: 4
  Filter action: Drop
  Nexthop type: none

  Level: 5
  Filter action: Drop
  Nexthop type: none

  Level: 6
  Filter action: Drop
  Nexthop type: none

  Level: 7
  Direction: down
  Filter action: Receive
  Nexthop type: none

Interface name: xe-0/0/0.0

  Level: 0
  Filter action: Drop
  Nexthop type: none

  Level: 1
  Filter action: Drop
  Nexthop type: none

  ...
```

show oam ethernet connectivity-fault-management forwarding-state interface interface-name

```
user@host> show oam ethernet connectivity-fault-management forwarding-state interface
interface-name ge-3/0/0.0
Interface name: ge-3/0/0.0

  Maintenance domain forwarding state:

  Level    Direction      Filter action     Nexthop      Nexthop
                                            type         index
  0                       Drop              none
  1                       Drop              none
  2                       Drop              none
  3                       Drop              none
```

```
4                       Drop            none
5                       Drop            none
6                       Drop            none
7       down            Receive         none
```

show oam ethernet connectivity-fault-management interfaces

Syntax **show oam ethernet connectivity-fault-management interfaces**
<*ethernet-interface-name*>
<level *md-level*>
<brief | detail | extensive>

Release Information Command introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management (CFM) database information for Ethernet interfaces.

Options brief | detail | extensive—(Optional) Display the specified level of output.

ethernet-interface-name—(Optional) Display CFM information only for CFM entities attached to the specified Ethernet interface.

level *md-level*—(Optional) Display CFM information for CFM identities enclosed within a maintenance domain of the specified level.

Required Privilege Level view

Related Documentation

- **clear oam ethernet connectivity-fault-management statistics on page 3493**
- **show oam ethernet connectivity-fault-management linktrace path-database on page 3504**
- **show oam ethernet connectivity-fault-management mep-database on page 3512**

List of Sample Output
show oam ethernet connectivity-fault-management interfaces on page 3501
show oam ethernet connectivity-fault-management interfaces detail on page 3501
show oam ethernet connectivity-fault-management interfacesextensive on page 3502
show oam ethernet connectivity-fault-management interfaces level on page 3503
show oam ethernet connectivity-fault-management interfaces (Trunk Interfaces) on page 3503

Output Fields Table 457 on page 3498 lists the output fields for the **show oam ethernet connectivity-fault-management interfaces** command. Output fields are listed in the approximate order in which they appear.

Table 457: show oam ethernet connectivity-fault-management interfaces Output Fields

Field Name	Field Description	Level of Output
Interface	Interface identifier.	All levels
Interface status	Local interface status.	All levels
Link status	Local link status. **Up**, **down**, or **oam-down**.	All levels

Table 457: show oam ethernet connectivity-fault-management interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
Maintenance domain name	Maintenance domain name.	**detail extensive**
Format (Maintenance domain)	Maintenance domain name format configured.	**detail extensive**
Level	Maintenance domain level configured.	All levels
Maintenance association name	Maintenance association name.	**detail extensive**
Format (Maintenance association)	Maintenance association name format configured.	**detail extensive**
Continuity-check status	Continuity-check status.	**detail extensive**
Interval	Continuity-check message interval.	**detail extensive**
Loss-threshold	Lost continuity-check message threshold.	**detail extensive**
MEP identifier	Maintenance association end point (MEP) identifier.	All levels
Neighbours	Number of MEP neighbors.	All levels
Direction	MEP direction configured.	**detail extensive**
MAC address	MAC address configured for the MEP.	**detail extensive**
MEP status	Indicates the status of the Connectivity Fault Management (CFM) protocol running on the MEP: **Running**, **inactive**, **disabled**, or **unsupported**.	**detail extensive**
Remote MEP not receiving CCM	Whether the remote MEP is not receiving connectivity check messages (CCMs).	**detail extensive**
Erroneous CCM received	Whether erroneous CCMs have been received.	**detail extensive**
Cross-connect CCM received	Whether cross-connect CCMs have been received.	**detail extensive**
RDI sent by some MEP	Whether the remote defect indication (RDI) bit is set in messages that have been received. The absence of the RDI bit in a CCM indicates that the transmitting MEP is receiving CCMs from all configured MEPs.	**detail extensive**
CCMs sent	Number of CCMs transmitted.	**detail extensive**

Table 457: show oam ethernet connectivity-fault-management interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
CCMs received out of sequence	Number of CCMs received out of sequence.	**detail extensive**
LBMs sent	Number of loopback request messages (LBMs) sent.	**detail extensive**
Valid in-order LBRs received	Number of loopback response messages (LBRs) received that were valid messages and in sequence.	**detail extensive**
Valid out-of-order LBRs received	Number of LBRs received that were valid messages and not in sequence.	**detail extensive**
LBRs received with corrupted data	Number of LBRs received that were corrupted.	**detail extensive**
LBRs sent	Number of LBRs transmitted.	**detail extensive**
LTMs sent	Linktrace messages (LTMs) transmitted.	**detail extensive**
LTMs received	Linktrace messages received.	**detail extensive**
LTRs sent	Linktrace responses (LTRs) transmitted.	**detail extensive**
LTRs received	Linktrace responses received.	**detail extensive**
Sequence number of next LTM request	Sequence number of next LTM request to be transmitted.	**detail extensive**
1DMs sent	If the interface is attached to an initiator MEP for a one-way ETH-DM session: Number of one-way delay measurement (1DM) PDU frames sent to the peer MEP in this session. For all other cases, this field displays **0**.	**detail extensive**
Valid 1DMs received	If the interface is attached to a receiver MEP for a one-way ETH-DM session: Number of valid 1DM frames received. For all other cases, this field displays **0**.	**detail extensive**
Invalid 1DMs received	If the interface is attached to a receiver MEP for a one-way ETH-DM session: Number of invalid 1DM frames received. For all other cases, this field displays **0**.	**detail extensive**
DMMs sent	If the interface is attached to an initiator MEP for a two-way ETH-DM session: Number of Delay Measurement Message (DMM) PDU frames sent to the peer MEP in this session. For all other cases, this field displays **0**.	**detail extensive**

Table 457: show oam ethernet connectivity-fault-management interfaces Output Fields *(continued)*

Field Name	Field Description	Level of Output
DMRs sent	If the interface is attached to a responder MEP for a two-way ETH-DM session: Number of Delay Measurement Reply (DMR) frames sent. For all other cases, this field displays **0**.	**detail extensive**
Valid DMRs received	If the interface is attached to an initiator MEP for a two-way ETH-DM session: Number of valid DMRs received. For all other cases, this field displays **0**.	**detail extensive**
Invalid DMRs received	If the interface is attached to an initiator MEP for a two-way ETH-DM session: Number of invalid DMRs received. For all other cases, this field displays **0**.	**detail extensive**
Remote MEP count	Number of remote MEPs.	**extensive**
Identifier (remote MEP)	MEP identifier of the remote MEP.	**extensive**
MAC address (remote MEP)	MAC address of the remote MEP.	**extensive**
State (remote MEP)	State of the remote MEP.	**extensive**
Interface (remote MEP)	Interface of the remote MEP.	**extensive**

show oam ethernet connectivity-fault-management interfaces

```
user@host> show oam ethernet connectivity-fault-management interfaces
Interface       Link       Status          Level       MEP          Neighbours
                                                       Identifier
ge-1/1/0.0      Up         Active          0           2            1
ge-1/1/0.1      Up         Active          0           2            1
ge-1/1/0.10     Up         Active          0           2            1
ge-1/1/0.100    Up         Active          0           2            1
ge-1/1/0.101    Up         Active          0           2            1
ge-1/1/0.102    Up         Active          0           2            1
ge-1/1/0.103    Up         Active          0           2            1
ge-1/1/0.104    Up         Active          0           2            1
ge-1/1/0.105    Up         Active          0           2            1
ge-1/1/0.106    Up         Active          0           2            1

...
```

show oam ethernet connectivity-fault-management interfaces detail

```
user@host> show oam ethernet connectivity-fault-management interfaces detail
Interface name: ge-5/2/9.0, Interface status: Active, Link status: Up
  Maintenance domain name: md0, Format: string, Level: 5
    Maintenance association name: ma1, Format: string
    Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
    MEP identifier: 1, Direction: down, MAC address: 00:90:69:0b:4b:94
```

```
  MEP status: running
  Defects:
    Remote MEP not receiving CCM                : no
    Erroneous CCM received                      : yes
    Cross-connect CCM received                  : no
    RDI sent by some MEP                        : yes
  Statistics:
    CCMs sent                                   : 76
    CCMs received out of sequence               : 0
    LBMs sent                                   : 0
    Valid in-order LBRs received                : 0
    Valid out-of-order LBRs received            : 0
    LBRs received with corrupted data           : 0
    LBRs sent                                   : 0
    LTMs sent                                   : 0
    LTMs received                               : 0
    LTRs sent                                   : 0
    LTRs received                               : 0
    Sequence number of next LTM request         : 0
    1DMs sent                                   : 0
    Valid 1DMs received                         : 0
    Invalid 1DMs received                       : 0
    DMMs sent                                   : 0
    DMRs sent                                   : 0
    Valid DMRs received                         : 0
    Invalid DMRs received                       : 0
  Remote MEP count: 2
    Identifier    MAC address        State    Interface
     2001     00:90:69:0b:7f:71       ok    ge-5/2/9.0
     4001     00:90:69:0b:09:c5       ok    ge-5/2/9.0
```

show oam ethernet connectivity-fault-management interfaces extensive

```
user@host> show oam ethernet connectivity-fault-management interfaces extensive
Interface name: ge-5/2/9.0, Interface status: Active, Link status: Up
  Maintenance domain name: md0, Format: string, Level: 5
    Maintenance association name: ma1, Format: string
    Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
    MEP identifier: 1, Direction: down, MAC address: 00:90:69:0b:4b:94
    MEP status: running
    Defects:
      Remote MEP not receiving CCM              : no
      Erroneous CCM received                    : yes
      Cross-connect CCM received                : no
      RDI sent by some MEP                      : yes
    Statistics:
      CCMs sent                                 : 76
      CCMs received out of sequence             : 0
      LBMs sent                                 : 0
      Valid in-order LBRs received              : 0
      Valid out-of-order LBRs received          : 0
      LBRs received with corrupted data         : 0
      LBRs sent                                 : 0
      LTMs sent                                 : 0
      LTMs received                             : 0
      LTRs sent                                 : 0
      LTRs received                             : 0
      Sequence number of next LTM request       : 0
      1DMs sent                                 : 0
      Valid 1DMs received                       : 0
      Invalid 1DMs received                     : 0
      DMMs sent                                 : 0
      DMRs sent                                 : 0
```

```
    Valid DMRs received                          : 0
    Invalid DMRs received                        : 0
  Remote MEP count: 2
    Identifier    MAC address        State    Interface
     2001     00:90:69:0b:7f:71       ok    ge-5/2/9.0
     4001     00:90:69:0b:09:c5       ok    ge-5/2/9.0
```

show oam ethernet connectivity-fault-management interfaces level

```
user@host> show oam ethernet connectivity-fault-management interfaces level 7
Interface       Link       Status           Level       MEP          Neighbours
                                                        Identifier
ge-3/0/0.0      Up         Active           7           201          0
xe-0/0/0.0      Up         Active           7           203          1
```

show oam ethernet connectivity-fault-management interfaces (Trunk Interfaces)

```
user@host> show oam ethernet connectivity-fault-management interfaces

Interface                   Link      Status       Level  MEP          Neighbours
                                                          Identifier
ge-4/0/1.0, vlan 100        Up        Active       5      100          0
ge-10/3/10.4091, vlan 4091  Down      Inactive     4      400          0
ge-4/0/0.0                  Up        Active       6      200          0

user@host> show oam ethernet connectivity-fault-management interfaces ge-4/0/0.0

Interface                   Link      Status       Level  MEP          Neighbours
                                                          Identifier
ge-4/0/0.0                  Up        Active       6      200          0

user@host> show oam ethernet connectivity-fault-management interfaces ge-4/0/1.0 vlan 100

Interface                   Link      Status       Level  MEP          Neighbours
                                                          Identifier
ge-4/0/1.0, vlan 100        Up        Active       5      100          0

user@host> show oam ethernet connectivity-fault-management interfaces ge-10/3/10.4091
vlan 4091

Interface                   Link      Status       Level  MEP          Neighbours
                                                          Identifier
ge-10/3/10.4091, vlan 4091  Down      Inactive     4      400          0
```

show oam ethernet connectivity-fault-management linktrace path-database

Syntax **show oam ethernet connectivity-fault-management path-database** *host* **maintenance-association** *ma-name* **maintenance-domain** *md-name mac-address*

Release Information Command introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management maintenance linktrace database information.

Options mac-address—Display connectivity fault management path database information for the specified MAC address of the remote host.

maintenance-association *ma-name*—Display connectivity fault management path database information for the specified maintenance association.

maintenance-domain *md-name*—Display connectivity fault management path database information for the specified maintenance domain.

Required Privilege Level view

Related Documentation

- **clear oam ethernet connectivity-fault-management statistics on page 3493**
- **show oam ethernet connectivity-fault-management interfaces on page 3498**
- **show oam ethernet connectivity-fault-management mip on page 3512**

List of Sample Output **show oam ethernet connectivity-fault-management path-database on page 3505**
show oam ethernet connectivity-fault-management linktrace path-database (Two traceroute Commands) on page 3505

Output Fields Table 458 on page 3504 lists the output fields for the **show oam ethernet connectivity-fault-management path-database** command. Output fields are listed in the approximate order in which they appear.

Table 458: show oam ethernet connectivity-fault-management linktrace path-database Output Fields

Field Name	Field Description
Linktrace to	MAC address of the 802.1ag node to which the linktrace message is targeted.
Interface	Interface used by the local MEP to send the linktrace message (LTM).
Maintenance Domain	Maintenance domain identifier specified in the traceroute command.
Maintenance Association	Maintenance association identifier specified in the traceroute command.
Level	Maintenance domain level configured for the maintenance domain.

Table 458: show oam ethernet connectivity-fault-management linktrace path-database Output Fields *(continued)*

Field Name	Field Description
Local Mep	MEP identifier of the local MEP originating the linktrace.
Hop	Sequential hop count of the linktrace path.
TTL	Number of hops remaining in the linktrace message (LTM). The time to live (TTL) is decremented at each hop.
Source MAC address	MAC address of the 802.1ag maintenance intermediate point (MIP) that is forwarding the LTM.
Next hop MAC address	MAC address of the 802.1ag node that is the next hop in the LTM path.
Transaction Identifier	4-byte identifier maintained by the MEP. Each LTM uses a transaction identifier. The transaction identifier is maintained globally across all maintenance domains. Use the transaction identifier to match an incoming linktrace responses (LTR), with a previously sent LTM.

show oam ethernet connectivity-fault-management path-database

```
user@host> show oam ethernet connectivity-fault-management path-database
maintenance-domain MD1 maintenance-association MA1 00:01:02:03:04:05
Linktrace to 00:01:02:03:04:05, Interface : ge-5/0/0.0
       Maintenance Domain: MD1, Level: 7
       Maintenance Association: MA1, Local Mep: 1

       Hop     TTL      Source MAC address        Next hop MAC address
       Transaction Identifier:100001
       1       63       00:00:aa:aa:aa:aa         00:00:bb:bb:bb:bb
       2       62       00:00:bb:bb:bb:bb         00:00:cc:cc:cc:cc
       3       61       00:00:cc:cc:cc:cc         00:01:02:03:04:05
       4       60       00:01:02:03:04:05         00:00:00:00:00:00
```

show oam ethernet connectivity-fault-management linktrace path-database (Two traceroute Commands)

```
user@host> show oam ethernet connectivity-fault-management path-database
maintenance-domain MD2 maintenance-association MA2 00:06:07:08:09:0A
Linktrace to 00:06:07:08:09:0A, Interface : ge-5/0/1.0
       Maintenance Domain: MD2, Level: 6
       Maintenance Association: MA2, Local Mep: 10

       Hop     TTL      Source MAC address        Next hop MAC address
       Transaction Identifier:100002
       1       63       00:00:aa:aa:aa:aa         00:00:bb:bb:bb:bb
       2       62       00:00:bb:bb:bb:bb         00:00:cc:cc:cc:cc
       3       61       00:00:cc:cc:cc:cc         00:06:07:08:09:0A
       4       60       00:06:07:08:09:0A         00:00:00:00:00:00
       Transaction Identifier:100003
       1       63       00:00:aa:aa:aa:aa         00:00:bb:bb:bb:bb
       2       62       00:00:bb:bb:bb:bb         00:00:cc:cc:cc:cc
       3       61       00:00:cc:cc:cc:cc         00:06:07:08:09:0A
       4       60       00:06:07:08:09:0A         00:00:00:00:00:00
```

show oam ethernet connectivity-fault-management mep-database

Syntax **show oam ethernet connectivity-fault-management mep-database**
maintenance-domain *domain-name*
maintenance-association *ma-name*
<local-mep *local-mep-id*>
<remote-mep *remote-mep-id*>

Release Information Command introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Display IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management (CFM) database information for CFM maintenance association end points (MEPs) in a CFM session.

Options
maintenance-association *ma-name*—Display connectivity fault management information for the specified maintenance association.

maintenance-domain *domain-name*—Display connectivity fault management information for the specified maintenance domain.

local-mep *local-mep-id*—(Optional) Display connectivity fault management information for the specified local MEP only.

remote-mep *remote-mep-id*—(Optional) Display connectivity fault management information for the specified remote MEP only.

Required Privilege Level view

Related Documentation
- **clear oam ethernet connectivity-fault-management statistics on page 3493**
- **show oam ethernet connectivity-fault-management interfaces on page 3498**
- **show oam ethernet connectivity-fault-management mip on page 3512**

List of Sample Output
show oam ethernet connectivity-fault-management mep-database on page 3510
show oam ethernet connectivity-fault-management mep-database local-mep remote-mep on page 3510
show oam ethernet connectivity-fault-management mep-database remote-mep (Action Profile Event) on page 3510

Output Fields Table 459 on page 3506 lists the output fields for the **show oam ethernet connectivity-fault-management mep-database** command. Output fields are listed in the approximate order in which they appear.

Table 459: show oam ethernet connectivity-fault-management mep-database Output Fields

Field Name	Field Description
Maintenance domain name	Maintenance domain name.

Table 459: show oam ethernet connectivity-fault-management mep-database Output Fields *(continued)*

Field Name	Field Description
Format (Maintenance domain)	Maintenance domain name format configured.
Level	Maintenance domain level configured.
Maintenance association name	Maintenance association name.
Format (Maintenance association)	Maintenance association name format configured.
Continuity-check status	Continuity-check status.
Interval	Continuity-check message interval.
MEP identifier	Maintenance association end point (MEP) identifier.
Direction	MEP direction configured.
MAC address	MAC address configured for the MEP.
Auto-discovery	Whether automatic discovery is enabled or disabled.
Priority	Priority used for CCMs and linktrace messages transmitted by the MEP.
Interface name	Interface identifier.
Interface status	Local interface status.
Link status	Local link status.
Remote MEP not receiving CCM	Whether the remote MEP is not receiving CCMs.
Erroneous CCM received	Whether erroneous CCMs have been received.
Cross-connect CCM received	Whether cross-connect CCMs have been received.
RDI sent by some MEP	Whether the remote defect indication (RDI) bit is set in messages that have been received. The absence of the RDI bit in a CCM indicates that the transmitting MEP is receiving CCMs from all configured MEPs.
CCMs sent	Number of CCMs transmitted.
CCMs received out of sequence	Number of CCMs received out of sequence.

Table 459: show oam ethernet connectivity-fault-management mep-database Output Fields *(continued)*

Field Name	Field Description
LBMs sent	Number of loopback messages (LBMs) sent.
Valid in-order LBRs received	Number of loopback response messages (LBRs) received that were valid messages and in sequence.
Valid out-of-order LBRs received	Number of LBRs received that were valid messages and not in sequence.
LBRs received with corrupted data	Number of LBRs received that were corrupted.
LBRs sent	Number of LBRs transmitted.
LTMs sent	Linktrace messages (LTMs) transmitted.
LTMs received	Linktrace messages received.
LTRs sent	Linktrace responses (LTRs) transmitted.
LTRs received	Linktrace responses received.
Sequence number of next LTM request	Sequence number of the next linktrace message request to be transmitted.
1DMs sent	If the MEP is an initiator for a one-way ETH-DM session: Number of one-way delay measurement (1DM) PDU frames sent to the peer MEP in this session. For all other cases, this field displays **0**.
Valid 1DMs received	If the MEP is a receiver for a one-way ETH-DM session: Number of valid 1DM frames received. For all other cases, this field displays **0**.
Invalid 1DMs received	If the MEP is a receiver for a one-way ETH-DM session: Number of invalid 1DM frames received. For all other cases, this field displays **0**.
DMMs sent	If the MEP is an initiator for a two-way ETH-DM session: Number of Delay Measurement Message (DMM) PDU frames sent to the peer MEP in this session. For all other cases, this field displays **0**.
DMRs sent	If the MEP is a responder for a ETH-DM session: Number of Delay Measurement Reply (DMR) frames sent. For all other cases, this field displays **0**.

Table 459: show oam ethernet connectivity-fault-management mep-database Output Fields *(continued)*

Field Name	Field Description
Valid DMRs received	If the MEP is an initiator for a two-way ETH-DM session: Number of valid DMRs received. For all other cases, this field displays **0**.
Invalid DMRs received	If the MEP is an initiator for a two-way ETH-DM session: Number of invalid DMRs received. For all other cases, this field displays **0**.
Remote MEP identifier	MEP identifier of the remote MEP.
State (remote MEP)	State of the remote MEP: **idle**, **start**, **ok**, or **failed**.
MAC address	MAC address of the remote MEP.
Type	Whether the remote MEP MAC address was learned using automatic discovery or configured.
Interface	Interface of the remote MEP. A seven-digit number is appended if CFM is configured to run on a routing instance of type VPLS.
Last flapped	Date, time, and how long ago the remote MEP interface went from down to up. The format is **Last flapped:** ***year-month-day hours*:*minutes*:*seconds timezone*** **(*hours*:*minutes*:*seconds* ago)**. For example, **Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago)**.
Remote defect indication	Whether the remote defect indication (RDI) bit is set in messages that have been received or transmitted.
Port status TLV	• In the Maintenance domain section, displays the last transmitted port status TLV value. • In the Remote MEP section, displays the last value of port status TLV received from the remote MEP. In the Action profile section, displays, the last occurred event **port-status-tlv blocked** event. This event occurred due to the reception of **blocked** value in the port status TLV from remote MEP.
Interface status TLV	• In the Maintenance domain section, displays the last transmitted interface status TLV value. • In the Remote MEP section, displays the last value of interface status TLV received from the remote MEP. In the Action profile section, if displays, the last occurred event interface-status-tlv event (either **lower-layer-down** or **down**). This event occurred due to the reception of either lower or **down** value in the interface status TLV from remote MEP.
Action profile	Name of the action profile occurrence associated with a remote MEP.
Last event	When an action profile occurs, displays the last event that triggered it.
Last event cleared	When all the configured and occurred events (under action profile) are cleared, then the action taken gets reverted (such as down interface is made up) and the corresponding time is noted and displayed.
Action	Action taken and the corresponding time of the action occurrence.

show oam ethernet connectivity-fault-management mep-database

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain vpls-vlan2000 maintenance-association vpls-vlan200
Maintenance domain name: vpls-vlan2000, Format: string, Level: 5
  Maintenance association name: vpls-vlan200, Format: string
  Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
  MEP identifier: 200, Direction: up, MAC address: 00:19:e2:b0:74:01
  Auto-discovery: enabled, Priority: 0
  Interface name: ge-0/0/1.0, Interface status: Active, Link status: Up
  Defects:
    Remote MEP not receiving CCM                 : no
    Erroneous CCM received                       : no
    Cross-connect CCM received                   : no
    RDI sent by some MEP                         : no
  Statistics:
    CCMs sent                                    : 1476
    CCMs received out of sequence                : 0
    LBMs sent                                    : 85
    Remote MEP count: 1
    Identifier    MAC address        State    Interface
      100     00:19:e2:b2:81:4b       ok  vt-0/1/10.1049088
```

show oam ethernet connectivity-fault-management mep-database local-mep remote-mep

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain vpls-vlan2000 maintenance-association vpls-vlan200 local-mep 200
remote-mep 100
Maintenance domain name: vpls-vlan2000, Format: string, Level: 5
  Maintenance association name: vpls-vlan200, Format: string
  Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
  MEP identifier: 200, Direction: up, MAC address: 00:19:e2:b0:74:01
  Auto-discovery: enabled, Priority: 0
  Interface name: ge-0/0/1.0, Interface status: Active, Link status: Up

  Remote MEP identifier: 100, State: ok
    MAC address: 00:19:e2:b2:81:4b, Type: Learned
    Interface: vt-0/1/10.1049088
    Last flapped: Never
    Remote defect indication: false
    Port status TLV: none
    Interface status TLV: none
```

show oam ethernet connectivity-fault-management mep-database remote-mep (Action Profile Event)

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md5 maintenance-association ma5 remote-mep 200
Maintenance domain name: md5, Format: string, Level: 5
  Maintenance association name: ma5, Format: string
  Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
  MEP identifier: 100, Direction: down, MAC address: 00:05:85:73:e8:ad
  Auto-discovery: enabled, Priority: 0
  Interface status TLV: none, Port status TLV: none
  Interface name: ge-1/0/8.0, Interface status: Active, Link status: Up

  Remote MEP identifier: 200, State: ok
    MAC address: 00:05:85:73:96:1f, Type: Configured
    Interface: ge-1/0/8.0
    Last flapped: Never
    Remote defect indication: false
    Port status TLV: none
    Interface status TLV: lower-layer-down
    Action profile: juniper
```

```
Last event: Interface-status-tlv lower-layer-down
Action: Interface-down, Time: 2009-03-27 14:25:10 PDT (00:00:02 ago)
```

show oam ethernet connectivity-fault-management mip

Syntax show oam ethernet connectivity-fault-management mip
<qualifier>

Release Information Command introduced in Junos OS Release 10.2 for J-EX Series switches.

Description Display all the maintenance association intermediate points (MIPs) created in the system. Specify qualifiers to display specific MIPs.

Options ***qualifier***—(Optional) Display the specified MIP.

Required Privilege Level view

Related Documentation

- **show oam ethernet connectivity-fault-management interfaces on page 3498**
- **show oam ethernet connectivity-fault-management linktrace path-database on page 3504**

List of Sample Output **show oam ethernet connectivity-fault-management mip on page 3512**

Output Fields Table 460 on page 3512 lists the output fields for the **show oam ethernet connectivity-fault-management mip** command. Output fields are listed in the approximate order in which they appear.

Table 460: show oam ethernet connectivity-fault-management mip Output Fields

Field Name	Field Description
MIP information for instance	Header for the MIP information showing the MIP name.
Interface	Interface type-dpc/pic/port.unit-number.
Level	MIP level configured.

show oam ethernet connectivity-fault -management mip

```
user@host> show oam ethernet connectivity-fault-management mip
MIP information for  __mip_name__

MIP information for  default-switch bd1

  Interface       Level
  ge-3/0/0.0          7
  ge-3/0/1.0          6
  ge-3/0/3.0          6
```

CHAPTER 132

Monitoring General Network Traffic and Hosts

Monitoring Hosts Using the J-Web Ping Host Tool

Purpose Use the J-Web ping host tool to verify that the host can be reached over the network. The output is useful for diagnosing host and network connectivity problems. The switch sends a series of ICMP echo (ping) requests to a specified host and receives ICMP echo responses.

Action To use the J-Web ping host tool:

1. Select **Troubleshoot>Ping Host**.
2. Next to Advanced options, click the expand icon.
3. Enter information into the Ping Host page, as described in Table 461 on page 3513.

   The Remote Host field is the only required field.
4. Click **Start**.

   The results of the ping operation are displayed in the main pane . If no options are specified, each ping response is in the following format:

   ```
   bytes bytes from ip-address: icmp_seq=number ttl=number time=time
   ```
5. To stop the ping operation before it is complete, click **OK**.

Meaning Table 461 on page 3513 lists the fields.

Table 461: J-Web Ping Host Field Summary

Field	Function	Your Action
Remote Host	Identifies the host to ping.	Type the hostname or IP address of the host to ping.
Advanced Options		

Table 461: J-Web Ping Host Field Summary *(continued)*

Field	Function	Your Action
Don't Resolve Addresses	Determines whether to display hostnames of the hops along the path.	• To suppress the display of the hop hostnames, select the check box. • To display the hop hostnames, clear the check box.
Interface	Specifies the interface on which the ping requests are sent.	Select the interface on which ping requests are sent from the list. If you select **any**, the ping requests are sent on all interfaces.
Count	Specifies the number of ping requests to send.	Select the number of ping requests to send from the list.
Don't Fragment	Specifies the Don't Fragment (DF) bit in the IP header of the ping request packet.	• To set the DF bit, select the check box. • To clear the DF bit, clear the check box.
Record Route	Sets the record route option in the IP header of the ping request packet. The path of the ping request packet is recorded within the packet and displayed in the main pane.	• To record and display the path of the packet, select the check box. • To suppress the recording and display of the path of the packet, clear the check box.
Type-of-Service	Specifies the type-of-service (TOS) value in the IP header of the ping request packet.	Select the decimal value of the TOS field from the list.
Routing Instance	Name of the routing instance for the ping attempt.	Select the routing instance name from the list.
Interval	Specifies the interval, in seconds, between transmissions of individual ping requests.	Select the interval from the list.
Packet Size	Specifies the size of the ping request packet.	Type the size, in bytes, of the packet. The size can be from 0 through 65468. The switch adds 8 bytes of ICMP header to the size.
Source Address	Specifies the source address of the ping request packet.	Type the source IP address.
Time-to-Live	Specifies the time-to-live (TTL) hop count for the ping request packet.	Select the TTL value from the list.
Bypass Routing	Determines whether ping requests are routed by means of the routing table. If the routing table is not used, ping requests are sent only to hosts on the interface specified in the Interface box. If the host is not on that interface, ping responses are not sent.	• To bypass the routing table and send the ping requests to hosts on the specified interface only, select the check box. • To route the ping requests using the routing table, clear the check box.

Related Documentation

- Monitoring Interface Status and Traffic on page 931

Monitoring Network Traffic Using Traceroute

Purpose Use the Traceroute page in the J-Web interface to trace a route between the switch and a remote host. You can use a traceroute task to display a list of waypoints between the switch and a specified destination host. The output is useful for diagnosing a point of failure in the path from the switch platform to the destination host and addressing network traffic latency and throughput problems.

Action To use the traceroute tool:

1. Select **Troubleshoot>Traceroute**.
2. Next to **Advanced options**, click the expand icon.
3. Enter information into the Traceroute page.

   The **Remote Host** field is the only required field.
4. Click **Start**.
5. To stop the traceroute operation before it is complete, click **OK** while the results of the traceroute operation are being displayed.

Meaning The switch generates the list of waypoints by sending a series of ICMP traceroute packets in which the time-to-live (TTL) value in the messages sent to each successive waypoint is incremented by 1. (The TTL value of the first traceroute packet is set to 1.) In this manner, each waypoint along the path to the destination host replies with a Time Exceeded packet from which the source IP address can be obtained.

The results of the traceroute operation are displayed in the main pane. If no options are specified, each line of the traceroute display is in the following format:

hop-number host (ip-address) [as-number] time1 time2 time3

The switch sends a total of three traceroute packets to each waypoint along the path and displays the round-trip time for each traceroute operation. If the switch times out before receiving a **Time Exceeded** message, an asterisk (*) is displayed for that round-trip time.

Table 462: Traceroute field summary

Field	Function	Your Action
Remote Host	Identifies the destination host of the traceroute.	Type the hostname or IP address of the destination host.
Advanced Options		
Don't Resolve Addresses	Determines whether hostnames of the hops along the path are displayed, in addition to IP addresses.	To suppress the display of the hop hostnames, select the check box.
Gateway	Specifies the IP address of the gateway to route through.	Type the gateway IP address.

Table 462: Traceroute field summary *(continued)*

Field	Function	Your Action
Source Address	Specifies the source address of the outgoing traceroute packets.	Type the source IP address.
Bypass Routing	Determines whether traceroute packets are routed by means of the routing table. If the routing table is not used, traceroute packets are sent only to hosts on the interface specified in the Interface box. If the host is not on that interface, traceroute responses are not sent.	To bypass the routing table and send the traceroute packets to hosts on the specified interface only, select the check box.
Interface	Specifies the interface on which the traceroute packets are sent.	From the list, select the interface on which traceroute packets are sent. If you select any, the traceroute requests are sent on all interfaces.
Time-to-live	Specifies the maximum time-to-live (TTL) hop count for the traceroute request packet.	From the list, select the TTL.
Type-of-Service	Specifies the type-of-service (TOS) value to include in the IP header of the traceroute request packet.	From the list, select the decimal value of the TOS field.
Resolve AS Numbers	Determines whether the autonomous system (AS) number of each intermediate hop between the router and the destination host is displayed.	To display the AS numbers, select the check box.

Related Documentation

- Connecting and Configuring a J-EX Series Switch (CLI Procedure) on page 161
- Connecting and Configuring a J-EX Series Switch (J-Web Procedure) on page 163
- Configuring Gigabit Ethernet Interfaces (J-Web Procedure) on page 909
- Monitoring Interface Status and Traffic on page 931

CHAPTER 133

Configuration Statements for General Network Management and Monitoring

archive-sites

Syntax

```
archive-sites {
    site-name;
}
```

Hierarchy Level [edit accounting-options file *filename*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Configure an archive site. If more than one site name is configured, an ordered list of archive sites for the accounting-data log files is created. When a file is archived, the router or switch attempts to transfer the file to the first URL in the list, moving to the next site only if the transfer does not succeed. The log file is stored at the archive site with a filename of the format ***router-name_log-filename_timestamp***.

Options ***site-name***—Any valid FTP URL to a destination. For information about specifying valid FTP URLs, see the *Junos OS System Basics Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.

Required Privilege Level snmp—To view this statement in the configuration.
snmp-control—To add this statement to the configuration.

Related Documentation

- Configuring Archive Sites

class-usage-profile

Syntax

```
class-usage-profile profile-name {
    file filename;
    interval minutes;
    source-classes {
        source-class-name;
    }
    destination-classes {
        destination-class-name;
    }
}
```

Hierarchy Level [edit accounting-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Create a class usage profile, which is used to log class usage statistics to a file in the **/var/log** directory. The class usage profile logs class usage statistics for the configured source classes on every interface that has **destination-class-usage** configured.

For information about configuring source classes, see the *Junos OS Routing Protocols Configuration Guide* at http://www.juniper.net/techpubs/software/junos/. For information about configuring source class usage, see the *Junos OS Network Interfaces Configuration Guide* at http://www.juniper.net/techpubs/software/junos/.

Options ***profile-name***—Name of the destination class profile.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Class Usage Profiles

counters

Syntax

```
counters {
  counter-name;
}
```

Hierarchy Level [edit accounting-options **filter-profile** *profile-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Names of counters for which filter profile statistics are collected. The packet and byte counts for the counters are logged to a file in the **/var/log** directory.

Options ***counter-name***—Name of the counter.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring the Counters

destination-classes

Syntax

```
destination-classes {
  destination-class-name;
}
```

Hierarchy Level [edit accounting-options **class-usage-profile** *profile-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the destination classes for which statistics are collected.

Options ***destination-class-name***—Name of the destination class to include in the source class usage profile.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring a Class Usage Profile

fields (for Interface Profiles)

Syntax

```
fields {
  field-name;
}
```

Hierarchy Level [edit accounting-options **interface-profile** *profile-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Statistics to collect in an accounting-data log file for an interface.

Options ***field-name***—Name of the field:

- **input-bytes**—Input bytes
- **input-errors**—Generic input error packets
- **input-multicast**—Input packets arriving by multicast
- **input-packets**—Input packets
- **input-unicast**—Input unicast packets
- **output-bytes**—Output bytes
- **output-errors**—Generic output error packets
- **output-multicast**—Output packets sent by multicast
- **output-packets**—Output packets
- **output-unicast**—Output unicast packets

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Interface Profile

file (Associating with a Profile)

Syntax file *filename*;

Hierarchy Level [edit accounting-options **class-usage-profile** *profile-name*],
[edit accounting-options **filter-profile** *profile-name*],
[edit accounting-options **interface-profile** *profile-name*],
[edit accounting-options **mib-profile** *profile-name*],
[edit accounting-options **routing-engine-profile** *profile-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the accounting log file associated with the profile.

Options ***filename***—Name of the log file. You must specify a filename already configured in the **file** statement at the **[edit accounting-options]** hierarchy level.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Interface Profile
- Configuring the Filter Profile
- Configuring the MIB Profile
- Configuring the Routing Engine Profile

file (Configuring a Log File)

Syntax

```
file filename {
  archive-sites {
    site-name;
  }
  files number;
  nonpersistent;
  size bytes;
  source-classes time;
  transfer-interval minutes;
}
```

Hierarchy Level [edit accounting-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify a log file to be used for accounting data.

Options ***filename***—Name of the file in which to write accounting data.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Accounting-Data Log Files

files

Syntax files *number*;

Hierarchy Level [edit accounting-options file *filename*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the maximum number of log files to be used for accounting data.

Options ***number***—The maximum number of files. When a log file (for example, **profilelog**) reaches its maximum size, it is renamed **profilelog.0**, then **profilelog.1**, and so on, until the maximum number of log files is reached. Then the oldest log file is overwritten. The minimum value for ***number*** is 3 and the default value is 10.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring Accounting-Data Log Files

filter-profile

Syntax

```
filter-profile profile-name {
  counters {
    counter-name;
  }
  file filename;
  interval minutes;
}
```

Hierarchy Level [edit accounting-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Create a profile to filter and collect packet and byte count statistics and write them to a file in the **/var/log** directory. To apply the profile to a firewall filter, you include the **accounting-profile** statement at the **[edit firewall filter *filter-name*]** hierarchy level. For more information about firewall filters, see the *Junos OS Network Interfaces Configuration Guide*.

Options ***profile-name***—Name of the filter profile.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Filter Profile

interface-profile

Syntax

```
interface-profile profile-name {
  fields {
    field-name;
  }
  file filename;
  interval minutes;
}
```

Hierarchy Level [edit accounting-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Create a profile to filter and collect error and packet statistics and write them to a file in the **/var/log** directory. You can specify an interface profile for either a physical or a logical interface.

Options ***profile-name***—Name of the interface profile.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Interface Profile

interval

Syntax interval *minutes*;

Hierarchy Level [edit accounting-options **class-usage-profile** *profile-name*],
[edit accounting-options **filter-profile** *profile-name*],
[edit accounting-options **interface-profile** *profile-name*],
[edit accounting-options **mib-profile** *profile-name*],
[edit accounting-options **routing-engine-profile** *profile-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify how often statistics are collected for the accounting profile.

Options ***minutes***—Length of time between each collection of statistics.

Range: 1 through 2880 minutes

Default: 30 minutes

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the Interface Profile
- Configuring the Filter Profile
- Configuring the MIB Profile
- Configuring the Routing Engine Profile

mib-profile

Syntax

```
mib-profile profile-name {
  file filename;
  interval minutes;
  object-names {
    mib-object-name;
  }
  operation operation-name;
}
```

Hierarchy Level [edit accounting-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Create a MIB profile to collect selected MIB statistics and write them to a file in the **/var/log** directory.

Options ***profile-name***—Name of the MIB statistics profile.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the MIB Profile

object-names

Syntax

```
object-names {
  mib-object-name;
}
```

Hierarchy Level [edit accounting-options **mib-profile** *profile-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the name of each MIB object for which MIB statistics are collected for an accounting-data log file.

Options ***mib-object-name***—Name of a MIB object. You can specify more than one MIB object name.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Configuring the MIB Profile

operation

Syntax operation *operation-name*;

Hierarchy Level [edit accounting-options **mib-profile** *profile-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the name of the operation used to collect MIB statistics for an accounting-data log file.

Options ***operation-name***—Name of the operation to use. You can specify a **get**, **get-next**, or **walk** operation.
Default: walk

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring the MIB Profile

routing-engine-profile

Syntax
```
routing-engine-profile profile-name {
  fields {
    field-name;
  }
  file filename;
  interval minutes;
}
```

Hierarchy Level [edit accounting-options]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Create a Routing Engine profile to collect selected Routing Engine statistics and write them to a file in the **/var/log** directory.

Options ***profile-name***—Name of the Routing Engine statistics profile.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring the Routing Engine Profile

size

Syntax size *bytes*;

Hierarchy Level [edit accounting-options **file** *filename*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify attributes of an accounting-data log file.

Options ***bytes***—Maximum size of each log file, in bytes, kilobytes (KB), megabytes (MB), or gigabytes (GB). When a log file (for example, **profilelog**) reaches its maximum size, it is renamed **profilelog.0**, then **profilelog.1**, and so on, until the maximum number of log files is reached. Then the oldest log file is overwritten. If you do not specify a size, the file is closed, archived, and renamed when the time specified for the transfer interval is exceeded.

Syntax: ***x*** to specify bytes, ***x***k to specify KB, ***x***m to specify MB, ***x***g to specify GB

Range: 256 KB through 1 GB

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring the Maximum Size of the File

source-classes

Syntax
```
source-classes {
  source-class-name;
}
```

Hierarchy Level [edit accounting-options **class-usage-profile** *profile-name*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the source classes for which statistics are collected.

Options ***source-class-name***—Name of the source class to include in the class usage profile.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring a Class Usage Profile

start-time

Syntax start-time *time*;

Hierarchy Level [edit accounting-options file *filename*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the start time for transfer of an accounting-data log file.

Options *time*—Start time for file transfer.
Syntax: *YYYY-MM-DD.hh:mm*

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring the Start Time for File Transfer

transfer-interval

Syntax transfer-interval *minutes*;

Hierarchy Level [edit accounting-options file *filename*]

Release Information Statement introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Specify the length of time the file remains open and receives new statistics before it is closed and transferred to an archive site.

Options *minutes*—Time the file remains open and receives new statistics before it is closed and transferred to an archive site.
Range: 5 through 2880 minutes
Default: 30 minutes

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Configuring the Transfer Interval of the File

CHAPTER 134

Operational Mode Commands for General Network Management and Monitoring

monitor traffic

Syntax

monitor traffic
<absolute-sequence>
<brief | detail | extensive>
<count *count*>
<interface *interface-name*>
<layer2-headers>
<matching *matching*>
<no-domain-names>
<no-promiscuous>
<no-resolve>
<no-timestamp>
<print-ascii>
<print-hex>
<resolve-timeout>
<size *size*>

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display packet headers or packets received and sent from the Routing Engine.

NOTE: Using the **monitor traffic** command can degrade router or switch performance.

Delays from DNS resolution can be eliminated by using the **no-resolve** option.

Options

none—(Optional) Display packet headers transmitted through fxp0. On a TX Matrix Plus router, display packet headers transmitted through **em0**.

brief | detail | extensive—(Optional) Display the specified level of output.

absolute-sequence—(Optional) Display absolute TCP sequence numbers.

count *count*—(Optional) Specify the number of packet headers to display (0 through 1,000,000). The monitor traffic command quits automatically after displaying the number of packets specified.

interface *interface-name*—(Optional) Specify the interface on which the **monitor traffic** command displays packet data. If no interface is specified, the **monitor traffic** command displays packet data arriving on the lowest-numbered interface.

layer2-headers—(Optional) Display the link-level header on each line.

matching *matching*—(Optional) Display packet headers that match a regular expression. Use matching expressions to define the level of detail with which the **monitor traffic** command filters and displays packet data.

no-domain-names—(Optional) Suppress the display of the domain portion of hostnames. With the **no-domain-names** option enabled, the **monitor traffic** command displays only team for the hostname **team.company.net**.

no-promiscuous—(Optional) Do not put the interface into promiscuous mode.

no-resolve—(Optional) Suppress reverse lookup of the IP addresses.

no-timestamp—(Optional) Suppress timestamps on displayed packets.

print-ascii—(Optional) Display each packet in ASCII format.

print-hex—(Optional) Display each packet, except the link-level header, in hexadecimal format.

resolve-timeout *timeout*—(Optional) Amount of time the router or switch waits for each reverse lookup before timing out. You can set the timeout for between 1 and 4,294,967,295 seconds. The default is 4 seconds. To display each packet, use the **print-ascii**, **print-hex**, or **extensive** option.

size *size*—(Optional) Read but not display up to the specified number of bytes for each packet. When set to **brief** output, the default packet size is 96 bytes and is adequate for capturing IP, ICMP, UDP, and TCP packet data. When set to **detail** and **extensive** output, the default packet size is 1514. The **monitor traffic** command truncates displayed packets if the matched data exceeds the configured size.

Additional Information

In the **monitor traffic** command, you can specify an expression to match by using the **matching** option and including the expression in quotation marks:

monitor traffic matching "*expression*"

Replace ***expression*** with one or more of the match conditions listed in Table 463 on page 3533.

Table 463: Match Conditions for the monitor traffic Command

Match Type	Condition	Description
Entity	**host** {*address* | *hostname*]	Matches packets that contain the specified address or hostname. The **host** match condition can be prepended with the protocol match conditions **arp**, **ip**, or **rarp**, or any of the directional match conditions.
	net *address*	Matches packets with source or destination addresses containing the specified network address.
	net *address***mask** *mask*	Matches packets containing the specified network address and subnet mask.
	port [*port-number* | *port-name*]	Matches packets containing the specified source or destination TCP or UDP port number or port name. In place of the numeric port address, you can specify a text synonym, such as **bgp** (179), **dhcp** (67), or **domain** (53) (the port numbers are also listed).

Table 463: Match Conditions for the monitor traffic Command *(continued)*

Match Type	Condition	Description
Directional	**dst**	Matches packets going to the specified destination. This match condition can be prepended to any of the entity type match conditions.
	src	Matches packets from a specified source. This match condition can be prepended to any of the entity type match conditions.
	src and **dst**	Matches packets that contain the specified source and destination addresses. This match condition can be prepended to any of the entity type match conditions.
	src or **dst**	Matches packets containing either of the specified addresses. This match condition can be prepended to any of the entity type match conditions.
Packet Length	**less** *value*	Matches packets shorter than or equal to the specified value, in bytes.
	greater *value*	Matches packets longer than or equal to the specified value, in bytes.
Protocol	**arp**	Matches all ARP packets.
	ether	Matches all Ethernet packets.
	ether [broadcast | multicast]	Matches broadcast or multicast Ethernet frames. This match condition can be prepended with**src** and **dst**.
	ether *protocol* **[***address* **| (arp | ip | rarp)]**	Matches packets with the specified Ethernet address or Ethernet packets of the specified protocol type. The **ether protocol** arguments **arp**, **ip**, and **rarp** are also independent match conditions, so they must be preceded by a backslash (\) when used in the **ether protocol** match condition.
	icmp	Matches all ICMP packets.
	ip	Matches all IP packets.
	ip [broadcast | multicast]	Matches broadcast or multicast IP packets.
	ip protocol [*address* **| (icmp | igrp | tcp | udp)]**	Matches packets with the specified address or protocol type. The **ip protocol** arguments **icmp**, **tcp**, and **udp** are also independent match conditions, so they must be preceded by a backslash (\) when used in the **ip protocol** match condition.
	isis	Matches all IS-IS routing messages.
	rarp	Matches all RARP packets.
	tcp	Matches all TCP datagrams.
	udp	Matches all UDP datagrams.

To combine expressions, use the logical operators listed in Table 464 on page 3535.

Table 464: Logical Operators for the monitor traffic Command

Logical Operator (Highest to Lowest Precedence)	Description
!	Logical NOT. If the first condition does not match, the next condition is evaluated.
&&	Logical AND. If the first condition matches, the next condition is evaluated. If the first condition does not match, the next condition is skipped.
||	Logical OR. If the first condition matches, the next condition is skipped. If the first condition does not match, the next condition is evaluated.
( )	Group operators to override default precedence order. Parentheses are special characters, each of which must be preceded by a backslash (\).

You can use relational operators to compare arithmetic expressions composed of integer constants, binary operators, a length operator, and special packet data accessors. The arithmetic expression matching condition uses the following syntax:

monitor traffic matching "ether[0] & 1 != 0"*"arithmetic_expression relational_operator arithmetic_expression"*

The packet data accessor uses the following syntax:

protocol [byte-offset <*size*>]

The optional *size* field represents the number of bytes examined in the packet header. The available values are **1, 2,** or **4** bytes.The following sample command captures all multicast traffic:

user@host> **monitor traffic matching "ether[0] & 1 != 0"**

To specify match conditions that have a numeric value, use the arithmetic and relational operators listed in Table 465 on page 3536.

NOTE: Because the Packet Forwarding Engine removes Layer 2 header information before sending packets to the Routing Engine:

- The monitor traffic command cannot apply match conditions to inbound traffic.
- The monitor traffic interface command also cannot apply match conditions for Layer 3 and Layer 4 packet data, resulting in the match pipe option (| match) for this command for Layer 3 and Layer 4 packets not working either. Therefore, ensure that you specify match conditions as described in this command summary. For more information about match conditions, see Table 463 on page 3533.
- The 802.1Q VLAN tag information included in the Layer 2 header is removed from all inbound traffic packets. As the monitor traffic interface ae[*x*] command for aggregated Ethernet interfaces (such as ae0) only shows inbound traffic data, the command does not show VLAN tag information in the output.

Table 465: Arithmetic and Relational Operators for the monitor traffic Command

Arithmetic or Relational Operator	Description
Arithmetic Operator	
+	Addition operator.
-	Subtraction operator.
/	Division operator.
&	Bitwise AND.
*	Bitwise exclusive OR.
|	Bitwise inclusive OR.
Relational Operator (Highest to Lowest Precedence)	
<=	If the first expression is less than or equal to the second, the packet matches.
>=	If the first expression is greater than or equal to the second, the packet matches.
<	If the first expression is less than the second, the packet matches.
>	If the first expression is greater than the second, the packet matches.
=	If the compared expressions are equal, the packet matches.

Table 465: Arithmetic and Relational Operators for the monitor traffic Command *(continued)*

Arithmetic or Relational Operator	Description
!=	If the compared expressions are unequal, the packet matches.

Required Privilege Level trace and maintenance

List of Sample Output

Output Fields When you enter this command, you are provided feedback on the status of your request.

monitor traffic count

```
user@host> monitor traffic count 2
listening on fxp0
04:35:49.814125  In my-server.home.net.1295 > my-server.work.net.telnet: . ack
4122529478 win 16798 (DF)
04:35:49.814185 Out my-server.work.net.telnet > my-server.home.net.1295: P
1:38(37) ack 0 win 17680 (DF) [tos 0x10]
```

monitor traffic detail count

```
user@host> monitor traffic detail count 2
listening on fxp0
04:38:16.265864  In my-server.home.net.1295 > my-server.work.net.telnet: . ack
4122529971 win 17678 (DF) (ttl 121, id 6812)
04:38:16.265926 Out my-server.work.net.telnet.telnet > my-server.home.net.1295:
P 1:38(37) ack 0 win 17680 (DF) [tos 0x10]  (ttl 6)
```

monitor traffic extensive (Absolute Sequence)

```
user@host> monitor traffic extensive no-domain-names no-resolve no-timestamp count 20
matching "tcp" absolute-sequence
listening on fxp0
 In 207.17.136.193.179 > 192.168.4.227.1024: . 4042780859:4042780859(0) ack
1845421797 win 16384 <nop,nop,timestamp 4935628 965951> [tos 0xc0]  (ttl )
 In 207.17.136.193.179 > 192.168.4.227.1024: P 4042780859:4042780912(53) ack
1845421797 win 16384 <nop,nop,timestamp 4935628 965951>: BGP [|BGP UPDAT)
 In 192.168.4.227.1024 > 207.17.136.193.179: P 1845421797:1845421852(55) ack
4042780912 win 16384 <nop,nop,timestamp 965951 4935628>: BGP [|BGP UPDAT)
...
```

monitor traffic extensive (Relative Sequence)

```
user@host> monitor traffic extensive no-domain-names no-resolve no-timestamp count 20
matching "tcp"
listening on fxp0
 In 172.24.248.221.1680 > 192.168.4.210.23: . 396159737:396159737(0) ack 1664980689
 win 17574 (DF) (ttl 121, id 50003)
Out 192.168.4.210.23 > 172.24.248.221.1680: P 1:40(39) ack 0 win 17680 (DF) [tos
 0x10]  (ttl 64, id 5394)
 In 207.17.136.193.179 > 192.168.4.227.1024: P 4042775817:4042775874(57) ack
```

```
1845416593 win 16384 <nop,nop,timestamp 4935379 965690>: BGP [|BGP UPDAT)
...
```

monitor traffic extensive count

```
user@host>  monitor traffic extensive count 5 no-domain-names no-resolve
listening on fxp013:18:17.406933  In 192.168.4.206.2723610880 > 172.17.28.8.2049:
 40 null (ttl 64, id 38367)13:18:17.407577  In 172.17.28.8.2049 >
192.168.4.206.2723610880: reply ok 28 null (ttl 61, id 35495)13:18:17.541140  In
 0:e0:1e:42:9c:e0 0:e0:1e:42:9c:e0 9000 60:                          0000 0100
0000 0000 0000 0000 0000 0000                       0000 0000 0000 0000 0000
0000 0000 0000                          0000 0000 0000 0000 0000 0000
000013:18:17.591513  In 172.24.248.156.4139 > 192.168.4.210.23: .
3556964918:3556964918(0) ack 295526518 win 17601 (DF) (ttl 121, id
14)13:18:17.591568 Out 192.168.4.210.23 > 172.24.248.156.4139: P 1:40(39) ack 0
win 17680 (DF) [tos 0x10]  (ttl 64, id 52376)
```

monitor traffic matching

```
user@host> monitor traffic matching "net 192.168.1.0/24"
verbose output suppressed, use <detail> or <extensive> for full protocol decode
Address resolution is ON. Use <no-resolve> to avoid any reverse lookup delay.
Address resolution timeout is 4s.
Listening on fxp0, capture size 96 bytes

Reverse lookup for 192.168.1.255 failed (check DNS reachability).
Other reverse lookup failures will not be reported.
Use no-resolve to avoid reverse lookups on IP addresses.

21:55:54.003511  In IP truncated-ip - 18 bytes missing! 192.168.1.17.netbios-ns
> 192.168.1.255.netbios-ns: UDP, length 50
21:55:54.003585 Out IP truncated-ip - 18 bytes missing! 192.168.1.17.netbios-ns
> 192.168.1.255.netbios-ns: UDP, length 50
21:55:54.003864  In arp who-has 192.168.1.17 tell 192.168.1.9
...
```

ping

Syntax

ping *host*
<bypass-routing>
<count *requests*>
<detail>
<do-not-fragment>
<inet | inet6>
<interface *source-interface*>
<interval *seconds*>
<logical-system (all | *logical-system-name*)>
<loose-source *value*>
<no-resolve>
<pattern *string*>
<rapid>
<record-route>
<routing-instance *routing-instance-name*>
<size *bytes*>
<source *source-address*>
<strict strict-source *value*>
<tos *type-of-service*>
<ttl *value*>
<verbose>
<wait *seconds*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Check host reachability and network connectivity. The **ping** command sends Internet Control Message Protocol (ICMP) ECHO_REQUEST messages to elicit ICMP ECHO_RESPONSE messages from the specified host. Type Ctrl+c to interrupt a ping command.

Options

host—IP address or hostname of the remote system to ping.

bypass-routing—(Optional) Bypass the normal routing tables and send ping requests directly to a system on an attached network. If the system is not on a directly attached network, an error is returned. Use this option to ping a local system through an interface that has no route through it.

count *requests*—(Optional) Number of ping requests to send. The range of values is **1** through **2,000,000,000**. The default value is an unlimited number of requests.

detail—(Optional) Include in the output the interface on which the ping reply was received.

do-not-fragment—(Optional) Set the do-not-fragment (DF) bit in the IP header of the ping packets.

inet—(Optional) Ping Packet Forwarding Engine IPv4 routes.

inet6—(Optional) Ping Packet Forwarding Engine IPv6 routes.

interface *source-interface*—(Optional) Interface to use to send the ping requests.

interval *seconds*—(Optional) How often to send ping requests. The range of values, in seconds, is **1** through infinity. The default value is **1**.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

loose-source *value*—(Optional) Intermediate loose source route entry (IPv4). Open a set of values.

no-resolve—(Optional) Do not attempt to determine the hostname that corresponds to the IP address.

pattern *string*—(Optional) Specify a hexadecimal fill pattern to include in the ping packet.

rapid—(Optional) Send ping requests rapidly. The results are reported in a single message, not in individual messages for each ping request. By default, five ping requests are sent before the results are reported. To change the number of requests, include the count option.

record-route—(Optional) Record and report the packet's path (IPv4).

routing-instance *routing-instance-name*—(Optional) Name of the routing instance for the ping attempt.

size *bytes*—(Optional) Size of ping request packets. The range of values, in bytes, is **0** through **65,468**. The default value is **56**, which is effectively 64 bytes because 8 bytes of ICMP header data are added to the packet.

source *source-address*—(Optional) IP address of the outgoing interface. This address is sent in the IP source address field of the ping request. If this option is not specified, the default address is usually the loopback interface (**lo.0**).

strict—(Optional) Use the strict source route option (IPv4).

strict-source *value*—(Optional) Intermediate strict source route entry (IPv4). Open a set of values.

tos *type-of-service*—(Optional) Set the type-of-service (ToS) field in the IP header of the ping packets. The range of values is **0** through **255**.

ttl *value*—(Optional) Time-to-live (TTL) value to include in the ping request (IPv6). The range of values is **0** through **255**.

verbose—(Optional) Display detailed output.

wait *seconds*—(Optional) Delay, in seconds, after sending the last packet. If this option is not specified, the default delay is **10** seconds. If this option is used without the count option, a default count of **5** packets is used.

Required Privilege Level network

List of Sample Output **ping hostname on page 3541**

ping hostname size count on page 3541
ping hostname rapid on page 3541

Output Fields When you enter this command, you are provided feedback on the status of your request. An exclamation point (**!**) indicates that an echo reply was received. A period (**.**) indicates that an echo reply was not received within the timeout period. An **x** indicates that an echo reply was received with an error code these packets are not counted in the received packets count. They are accounted for separately.

ping hostname

```
user@host> ping skye
PING skye.net (192.168.169.254): 56 data bytes
64 bytes from 192.168.169.254: icmp_seq=0 ttl=253 time=1.028 ms
64 bytes from 192.168.169.254: icmp_seq=1 ttl=253 time=1.053 ms
64 bytes from 192.168.169.254: icmp_seq=2 ttl=253 time=1.025 ms
64 bytes from 192.168.169.254: icmp_seq=3 ttl=253 time=1.098 ms
64 bytes from 192.168.169.254: icmp_seq=4 ttl=253 time=1.032 ms
64 bytes from 192.168.169.254: icmp_seq=5 ttl=253 time=1.044 ms
^C [abort]
```

ping hostname size count

```
user@host> ping skye size 200 count 5
PING skye.net (192.168.169.254): 200 data bytes
208 bytes from 192.168.169.254: icmp_seq=0 ttl=253 time=1.759 ms
208 bytes from 192.168.169.254: icmp_seq=1 ttl=253 time=2.075 ms
208 bytes from 192.168.169.254: icmp_seq=2 ttl=253 time=1.843 ms
208 bytes from 192.168.169.254: icmp_seq=3 ttl=253 time=1.803 ms
208 bytes from 192.168.169.254: icmp_seq=4 ttl=253 time=17.898 ms

--- skye.net ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 1.759/5.075/17.898 ms
```

ping hostname rapid

```
user@host> ping skye rapid
PING skye.net (192.168.169.254): 56 data bytes
!!!!!
--- skye.net ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.956/0.974/1.025/0.026 ms
```

show snmp mib

Syntax show snmp mib (get | get-next | walk) (ascii | decimal) *object-id* .

Release Information Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description Display local Simple Network Management Protocol (SNMP) Management Information Base (MIB) object values.

Options get—Retrieve and display one or more SNMP object values.

get-next—Retrieve and display the next SNMP object values.

walk—Retrieve and display the SNMP object values that are associated with the requested object identifier (OID). When you use this option, the Junos OS displays the objects below the subtree that you specify.

ascii—Display the SNMP object's string indices as an ascii-key representation.

decimal—Display the SNMP object values in the decimal (default) format. The **decimal** option is the default option for this command. Therefore, issuing the **show snmp mib (get | get-next | walk) decimal object-id** and the **show snmp mib (get | get-next | walk) object-id** commands display the same output.

object-id—The object can be represented by a sequence of dotted integers (such as **1.3.6.1.2.1.2**) or by its subtree name (such as **interfaces**). When entering multiple objects, enclose the objects in quotation marks.

Required Privilege Level snmp—To view this statement in the configuration.

List of Sample Output

show snmp mib get on page 3543
show snmp mib get (Multiple Objects) on page 3543
show snmp mib get-next on page 3543
show snmp mib get-next (Specify an OID) on page 3543
show snmp mib walk on page 3543
show snmp mib walk decimal on page 3543
show snmp mib walk (ASCII) on page 3543
show snmp mib walk (Multiple Indices) on page 3543
show snmp mib walk decimal (Multiple Indices) on page 3543

Output Fields Table 466 on page 3542 describes the output fields for the **show snmp mib** command. Output fields are listed in the approximate order in which they appear.

Table 466: show snmp mib Output Fields

Field Name	Field Description
name	Object name and numeric instance value.
object value	Object value. The Junos OS translates OIDs into the corresponding object names.

show snmp mib get

```
user@host> show snmp mib get sysObjectID.0
sysObjectID.0 = jnxProductNameM20
```

show snmp mib get (Multiple Objects)

```
user@host> show snmp mib get ?sysObjectID.0 sysUpTime.0?
sysObjectID.0 = jnxProductNameM20
sysUpTime.0 = 1640992
```

show snmp mib get-next

```
user@host> show snmp mib get-next jnxMibs
jnxBoxClass.0 = jnxProductLineM20.0
```

show snmp mib get-next (Specify an OID)

```
user@host> show snmp mib get-next 1.3.6.1
sysDescr.0    = Juniper Networks, Inc. m20 internet router, kernel
JUNOS release: 2004-1 Build date: build date UTC Copyright (c) 1996-2004 Juniper
 Networks, Inc.
```

show snmp mib walk

```
user@host> show snmp mib walk system
sysDescr.0    = Juniper Networks, Inc. m20 internet router, kernel
JUNOS release #0: 2004-1 Build date: build date UTC Copyright (c) 1996-2004 Juniper
 Networks, Inc.
sysObjectID.0 = jnxProductNameM20
sysUpTime.0 = 1640992
sysContact.0 = Your contact
sysName.0 = my router
sysLocation.0 = building 1
sysServices.0 = 4
```

show snmp mib walk decimal

```
user@host show snmp mib walk decimal jnxUtilData
jnxUtilCounter32Value.102.114.101.100 = 100
```

show snmp mib walk (ASCII)

```
show snmp mib walk ascii jnxUtilData
jnxUtilCounter32Value."fred" = 100
```

show snmp mib walk (Multiple Indices)

```
show snmp mib walk ascii jnxFWCounterByteCount
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_BE-fe-1/3/0.0-i".2 = 0
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_CC-fe-1/3/0.0-i".2 = 0
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_RT-fe-1/3/0.0-i".2 = 0
.......
```

show snmp mib walk decimal (Multiple Indices)

```
show snmp mib walk ascii jnxFWCounterByteCount
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_BE-fe-1/3/0.0-i".2 = 0
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_CC-fe-1/3/0.0-i".2 = 0
jnxFWCounterByteCount."fe-1/3/0.0-i"."CLASS_RT-fe-1/3/0.0-i".2 = 0
.......
```

traceroute

Syntax

traceroute *host*
<as-number-lookup>
<bypass-routing>
<clns>
<gateway *address*>
<inet | inet6>
<interface *interface-name*>
<logical system (all | *logical-system-name*)>
<mpls (ldp *FEC address* | rsvp *label-switched-path-name*)>
<no-resolve>
<routing-instance *routing-instance-name*>
<source *source-address*>
<tos *value*>
<ttl *value*>
<wait *seconds*>

Release Information

Command introduced before Junos OS Release 10.2 for J-EX Series switches.

Description

Display the route packets take to a specified network host. Use **traceroute** as a debugging tool to locate points of failure in a network.

Options

host—IP address or name of remote host.

as-number-lookup—(Optional) Display the autonomous system (AS) number of each intermediate hop on the path from the host to the destination.

bypass-routing—(Optional) Bypass the normal routing tables and send requests directly to a system on an attached network. If the system is not on a directly attached network, an error is returned. Use this option to display a route to a local system through an interface that has no route through it.

clns—(Optional) Trace the route belonging to Connectionless Network Service (CLNS).

gateway *address*—(Optional) Address of a router or switch through which the route transits.

inet | inet6—(Optional) Trace the route belonging to IPv4 or IPv6, respectively.

interface *interface-name*—(Optional) Name of the interface over which to send packets.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

mpls (ldp *FEC address* | rsvp *label-switched-path name*)—(Optional) Analyze the status of LDP-signaled or RSVP-signaled MPLS label-switched paths (LSPs). You can optionally specify the forward equivalence class (FEC) address for the LDP LSP or the LSP name for RSVP. You can also analyze a specific LSP by issuing the **traceroute mpls rsvp *lsp-name*** command. You can only analyze IPv4 point-to-point LSPs. IPv6 is not supported.

no-resolve—(Optional) Do not attempt to determine the hostname that corresponds to the IP address.

routing-instance *routing-instance-name*—(Optional) Name of the routing instance for the traceroute attempt.

source *source-address*—(Optional) Source address of the outgoing traceroute packets.

tos *value*—(Optional) Value to include in the IP type-of-service (ToS) field. The range of values is **0** through **255**.

ttl *value*—(Optional) Maximum time-to-live value to include in the traceroute request. The range of values is **0** through **128**.

wait *seconds*—(Optional) Maximum time to wait for a response to the traceroute request.

Required Privilege Level network

List of Sample Output

traceroute on page 3545
traceroute as-number-lookup host on page 3546
traceroute noresolve on page 3546
traceroute (Between CE Routers, Layer 3 VPN) on page 3546
traceroute (Through an MPLS LSP) on page 3546

Output Fields Table 467 on page 3545 describes the output fields for the **traceroute** command. Output fields are listed in the approximate order in which they appear.

Table 467: traceroute Output Fields

Field Name	Field Description
traceroute to	IP address of the receiver.
hops max	Maximum number of hops allowed.
byte packets	Size of packets being sent.
number-of-hops	Number of hops from the source to the named router or switch.
router-name	Name of the router or switch for this hop.
address	Address of the router or switch for this hop.
Round trip time	Average round-trip time, in milliseconds (ms).

traceroute

```
user@host> traceroute santacruz
traceroute to green.company.net (10.156.169.254), 30 hops max, 40 byte packets
 1  blue23 (10.168.1.254)    2.370 ms  2.853 ms  0.367 ms
 2  red14  (10.168.255.250)  0.778 ms  2.937 ms  0.446 ms
```

```
 3  yellow (10.156.169.254)  7.737 ms  89.905 ms  0.834 ms
```

traceroute as-number-lookup host

```
user@host> traceroute as-number-lookup 10.100.1.1
traceroute to 10.100.1.1 (10.100.1.1), 30 hops max, 40 byte packets
 1  10.39.1.1 (10.39.1.1)  0.779 ms  0.728 ms  0.562 ms
 2  10.39.1.6 (10.39.1.6) [AS  32]  0.657 ms  0.611 ms  0.617 ms
 3  10.100.1.1 (10.100.1.1) [AS  10, 40, 50]  0.880 ms  0.808 ms  0.774 ms
```

traceroute noresolve

```
user@host> traceroute santacruz noresolve
traceroute to green.company.net (10.156.169.254), 30 hops max, 40 byte packets
 1  10.168.1.254  0.458 ms  0.370 ms  0.365 ms
 2  10.168.255.250  0.474 ms  0.450 ms  0.444 ms
 3  10.156.169.254  0.931 ms  0.876 ms  0.862 ms
```

traceroute (Between CE Routers, Layer 3 VPN)

```
user@host> traceroute vpn09
traceroute to vpn09.skybank.net (10.255.14.179), 30 hops max, 40
      byte packets
       1  10.39.10.21 (10.39.10.21)  0.598 ms  0.500 ms  0.461 ms
       2  10.39.1.13 (10.39.1.13)  0.796 ms  0.775 ms  0.806 ms
           MPLS Label=100006 CoS=0 TTL=1 S=1
       3  vpn09.skybank.net (10.255.14.179)  0.783 ms  0.716 ms  0.686
```

traceroute (Through an MPLS LSP)

```
user@host> traceroute mpls1
traceroute to 10.168.1.224 (10.168.1.224), 30 hops max, 40 byte packets
 1  mpls1-sr0.company.net (10.168.200.101)  0.555 ms  0.393 ms  0.367 ms
     MPLS Label=1024 CoS=0 TTL=1
 2  mpls5-lo0.company.net (10.168.1.224)  0.420 ms  0.394 ms  0.401 ms
```

PART 25

Index

- Index on page 3549

Index

Symbols

A

C

D

E

F

G

H

J

K

L

M

N

O

Q

R

S

T

U

V

W

X

Y